

Big Learning with Bayesian Methods

Jun Zhu, *Member, IEEE*, Jianfei Chen, Wenbo Hu

Abstract—Explosive growth in data and availability of cheap computing resources have sparked increasing interest in Big learning, an emerging subfield that studies scalable machine learning algorithms, systems, and applications with Big Data. Bayesian methods represent one important class of statistic methods for machine learning, with substantial recent developments on adaptive, flexible and scalable Bayesian learning. This article provides a survey of the recent advances in Big learning with Bayesian methods, termed Big Bayesian Learning, including nonparametric Bayesian methods for adaptively inferring model complexity, regularized Bayesian inference for improving the flexibility via posterior regularization, and scalable algorithms and systems based on stochastic subsampling and distributed computing for dealing with large-scale applications.

Index Terms—Big Bayesian Learning, Bayesian nonparametrics, Regularized Bayesian inference, Scalable algorithms

1 INTRODUCTION

WE live in an era of Big Data, where science, engineering and technology are producing massive data streams, with petabyte and exabyte scales becoming increasingly common [45], [57], [154]. Besides the explosive growth in volume, Big Data also has high velocity, high variety, and high uncertainty. These complex data streams require ever-increasing processing speeds, economical storage, and timely response for decision making in highly uncertain environments.

With the primary goal of building intelligent systems that automatically improve from experiences, machine learning (ML) is becoming an increasingly important field to tackle the big data challenges [128], with an emerging field of *Big Learning*, which covers theories, algorithms and systems on addressing big data problems.

1.1 Big Learning Challenges

In big data era, machine learning needs to deal with the challenges of learning from complex situations with *large N*, *large P*, *large L*, and *large M*, where *N* is the data size, *P* is the feature dimension, *L* is the number of tasks, and *M* is the model size. Given that *N* is obvious, we explain the other factors below.

Large P: with the development of Internet, data sets with ultrahigh dimensionality have emerged, such as the spam filtering data with trillion features [189] and the even higher-dimensional feature space via explicit kernel mapping [173]. Note that whether a learning problem is high-dimensional depends on the ratio between *P* and *N*. Many scientific problems with $P \gg N$ impose great challenges on learning, calling for effective regularization techniques to avoid overfitting and select salient features [64].

Large L: many tasks involve classifying text or images into tens of thousands or millions of categories. For example, the ImageNet [5] database consists of more than 14 millions of web images from 21 thousands of concepts, while with the goal of providing on average 1,000 images for each of 100+ thousands of concepts (or synsets) in WordNet; and the LSHTC text classification challenge 2014 aims to classify Wikipedia documents into one of 325,056 categories [2]. Often, these categories are organized in a graph, e.g., the tree structure in ImageNet and the DAG (directed acyclic graph) structure in LSHTC, which can be explored for better learning [28], [56].

Large M: with the availability of massive data, models with millions or billions of parameters are becoming common. Significant progress has been made on learning deep models, which have multiple layers of non-linearities allowing them to extract multi-grained representations of data, with successful applications in computer vision, speech recognition, and natural language processing. Such models include neural networks [84], auto-encoders [182], [108], and probabilistic generative models [161], [155].

1.2 Big Bayesian Learning

Though Bayesian methods have been widely used in ML, skepticism often arises when we talking about Bayesian methods for big data [94]. Practitioners also criticize that Bayesian methods are often too slow for even small-scaled problems, owing to many factors such as the non-conjugacy models with intractable integrals. Nevertheless, Bayesian methods have several advantages on dealing with:

- 1) **Uncertainty:** our world is an uncertain place because of physical randomness, incomplete knowledge, ambiguities and contradictions. Bayesian methods provide a principled theory for combining prior knowledge and uncertain evidence.

• J. Zhu, J. Chen and W. Hu are with TNList; State Key IST Lab; Dept. of Comp. Sci. & Tech.; Tsinghua University, Beijing, 100084 China. Email: {dcszj@, chenjian14@mails., hwb3@mails.}tsinghua.edu.cn

- 2) **Flexibility:** Bayesian methods are intuitively simple and flexible. Hierarchical Bayesian modeling offers a flexible tools for characterizing uncertainty, missing values, latent structures, and more.
- 3) **Adaptivity:** The dynamics and uncertainty of Big Data require that our models should be adaptive when the learning scenarios change. Nonparametric Bayesian methods provide elegant tools to deal with situations in which phenomena continue to emerge as data are collected [85]. Moreover, the Bayesian updating rule and its variants are sequential in nature and suitable for dealing with big data streams.
- 4) **Overfitting:** Although the data volume grows exponentially, the predictive information grows slower than the amount of Shannon information [30], while our models are becoming increasingly large by leveraging powerful computers, such as the deep networks with billions of parameters. It implies that our models are increasing their capacity faster than the amount of information that we need to fill them with, therefore causing serious overfitting problems that call for effective regularization [170].

Therefore, Bayesian methods are becoming increasingly relevant in the big data era [190] to protect high capacity models against overfitting, and to allow models adaptively updating their capacity. However, the application of Bayesian methods to big data problems runs into a computational bottleneck that needs to be addressed with new (approximate) inference methods. This article aims to provide a literature survey of the recent advances in big learning with Bayesian methods, including nonparametric Bayesian methods, regularized Bayesian inference, scalable inference algorithms and systems based on stochastic subsampling and distributed computing.

2 BASICS OF BAYESIAN METHODS

The general blueprint of Bayesian data analysis [67] is that a Bayesian model expresses a generative process of the data that includes hidden variables, under some statistical assumptions. The process specifies a joint probability distribution of the hidden and observed random variables. Given a set of observed data, data analysis is performed by *posterior inference*, which computes the conditional distribution of the hidden variables given the observed data.

2.1 Bayes' Theorem

At the core of Bayesian methods is Bayes' theorem (a.k.a Bayes' rule). Let Θ be the model parameters and \mathcal{D} be the given data set. The Bayesian posterior distribution is

$$p(\Theta|\mathcal{D}) = \frac{p_0(\Theta)p(\mathcal{D}|\Theta)}{p(\mathcal{D})}, \quad (1)$$

where $p_0(\cdot)$ is a prior distribution, chosen before seeing any data; $p(\mathcal{D}|\Theta)$ is the assumed likelihood model; and $p(\mathcal{D}) = \int p_0(\Theta)p(\mathcal{D}|\Theta)d\Theta$ is the marginal likelihood (or evidence), often involving an intractable integration problem that requires approximate inference as detailed below. The year 2013 marks the 250th anniversary of Thomas Bayes' essay on how humans can sequentially learn from experience, steadily updating their beliefs as more data become available [63].

A useful variational formulation of Bayes' rule is

$$\min_{q(\Theta) \in \mathcal{P}} \text{KL}(q(\Theta)||p_0(\Theta)) - \mathbb{E}_q[\log p(\mathcal{D}|\Theta)], \quad (2)$$

where \mathcal{P} is the space of all distributions that make the objective well-defined. It can be shown that the optimum solution to (2) is identical to the Bayesian posterior. In fact, if we add the constant term $\log p(\mathcal{D})$, the problem is equivalent to minimizing the KL-divergence between $q(\Theta)$ and the Bayesian posterior $p(\Theta|\mathcal{D})$, which is non-negative and takes 0 if and only if q equals to $p(\Theta|\mathcal{D})$. The variational interpretation is significant in two aspects: (1) it provides a basis for variational Bayes methods; and (2) it provides a starting point to make Bayesian methods more flexible by incorporating a rich set of posterior constraints. We will make these clear soon later.

It is noteworthy that $q(\Theta)$ represents the density of a general post-data posterior in the sense of [75, pp.15], not necessarily corresponding to a Bayesian posterior induced by Bayes' rule. As we shall see in Section 3.2, when we introduce additional constraints, the post-data posterior $q(\Theta)$ is different from the Bayesian posterior $p(\Theta|\mathcal{D})$, and moreover, it could even not be obtainable by the conventional Bayesian inference via Bayes' rule. In the sequel, in order to distinguish $q(\cdot)$ from the Bayesian posterior, we will call it post-data posterior. The optimization formulation in (2) implies that Bayes' rule is an information projection procedure that projects a prior density to a post-data posterior by taking account of the observed data. In general, Bayes's rule is a special case of the principle of minimum information [193].

2.2 Bayesian Methods in Machine Learning

Bayesian statistics has been applied to almost every ML task, ranging from the single-variate regression/classification to the structured output predictions and to the unsupervised/semi-supervised learning scenarios [31]. In essence, however, there are several basic tasks that we briefly review below.

Prediction: After training, Bayesian models make predictions using the distribution:

$$p(\mathbf{x}|\mathcal{D}) = \int p(\mathbf{x}, \Theta|\mathcal{D})d\Theta = \int p(\mathbf{x}|\Theta, \mathcal{D})p(\Theta|\mathcal{D})d\Theta, \quad (3)$$

where $p(\mathbf{x}|\Theta, \mathcal{D})$ is often simplified as $p(\mathbf{x}|\Theta)$ due to the i.i.d assumption of the data when the model is

given. Since the integral is taken over the posterior distribution, the training data is considered.

Model Selection: Model selection is a fundamental problem in statistics and machine learning [96]. Let \mathbf{M} be a family of models, where each model is indexed by a set of parameters Θ . Then, the marginal likelihood of the model family (or model evidence) is

$$p(\mathcal{D}|\mathbf{M}) = \int p(\mathcal{D}|\Theta)p(\Theta|\mathbf{M})d\Theta, \quad (4)$$

where $p(\Theta|\mathbf{M})$ is often assumed to be uniform if no strong prior exists.

For two different model families \mathbf{M}_1 and \mathbf{M}_2 , the ratio of model evidences $\kappa = \frac{p(\mathcal{D}|\mathbf{M}_1)}{p(\mathcal{D}|\mathbf{M}_2)}$ is called Bayes factor [98]. The advantage of using Bayes factors for model selection is that it automatically and naturally includes a penalty for including too much model structure [31, Chap 3]. Thus, it guards against overfitting. For models where an explicit version of the likelihood is not available or too costly to evaluate, approximate Bayesian computation (ABC) can be used for model selection in a Bayesian framework [79], [180], while with the caveat that approximate-Bayesian estimates of Bayes factors are often biased [157].

2.3 Approximate Bayesian Inference

The computational difficulties in Bayesian inference arise from the intractability of high-dimensional integrals as involved in the posterior and in Eqs (3, 4). These are typically not only analytically intractable but also difficult to obtain numerically. Common practice resorts to approximate methods, which can be grouped into two categories¹ — variational methods and Monte Carlo methods.

2.3.1 Variational Bayesian Methods

Variational methods have a long history in physics, statistics, control theory and economics. In machine learning, variational formulations appear naturally in regularization theory, maximum entropy estimates, and approximate inference in graphical models. We refer the readers to the seminal book [183] and the nice short overview [95] for more details. A variational method basically consists of two parts:

- 1) cast the problems as some optimization problems;
- 2) find an approximate solution when the exact solution is not feasible.

For Bayes' rule, we have provided a variational formulation in (2), which is equivalent to minimizing the KL-divergence between the variational distribution $q(\Theta)$ and the target posterior $p(\Theta|\mathcal{D})$. We can also

1. Both maximum likelihood estimation (MLE), $\hat{\Theta}_{\text{MLE}} = \arg\max_{\Theta} p(\mathcal{D}|\Theta)$, and maximum a posterior estimation (MAP), $\hat{\Theta}_{\text{MAP}} = \arg\max_{\Theta} p_0(\Theta)p(\mathcal{D}|\Theta)$, can be seen as the third type of approximation methods to do Bayesian inference. We omit them since they examine only a single point, and so can neglect the potentially large distributions in the integrals.

show that the negative of the objective in (2) is a lower bound of the evidence (i.e., log-likelihood):

$$\log p(\mathcal{D}) \geq \mathbb{E}_q[\log p(\Theta, \mathcal{D})] - \mathbb{E}_q[\log q(\Theta)]. \quad (5)$$

Then, variational Bayesian methods maximize the Evidence Lower BOund (ELBO):

$$\max_{q \in \mathcal{P}} \mathbb{E}_q[\log p(\Theta, \mathcal{D})] - \mathbb{E}_q[\log q(\Theta)], \quad (6)$$

whose solution is the target posterior if no assumptions are made.

However, in many cases it is intractable to calculate the target posterior. Therefore, to simplify the optimization, the variational distribution is often assumed to be in some parametric family, e.g., $q_{\phi}(\Theta)$, and has some mean-field representation

$$q_{\phi}(\Theta) = \prod_i q_{\phi_i}(\Theta_i), \quad (7)$$

where $\{\Theta_i\}$ represent a partition of Θ . Then, the problem transforms to find the best parameters $\hat{\phi}$ that maximize the ELBO, which can be solved with numerical optimization methods. For example, with the factorization assumption, coordinate descent is often used to iteratively solve for ϕ_i until reaching some local optimum. Once a variational approximation q^* is found, the Bayesian integrals can be approximated by replacing $p(\Theta|\mathcal{D})$ by q^* . In many cases, the model Θ consists of parameters θ and hidden variables \mathbf{h} . Then, if we make the (structured) mean-field assumption that $q(\theta, \mathbf{h}) = q(\theta)q(\mathbf{h})$, the variational problem can be solved by a variational Bayesian EM algorithm [24], which alternately updates $q(\mathbf{h})$ at the variational Bayesian E-step and updates $q(\theta)$ at the variational Bayesian M-step.

2.3.2 Monte Carlo Methods

Monte Carlo (MC) methods represent a diverse class of algorithms that rely on repeated random sampling to compute the solution to problems whose solution space is too large to explore systematically or whose systemic behavior is too complex to model. The basic idea of MC methods is to draw a set of i.i.d samples $\{\Theta_i\}_{i=1}^N$ from a target distribution $p(\Theta)$ and use the empirical distribution $\hat{p}(\cdot) = \frac{1}{N} \sum_{i=1}^N \delta_{\Theta_i}(\cdot)$, to approximate the target distribution, where $\delta_{\Theta_i}(\cdot)$ is the delta-Dirac mass located at Θ_i . Consider the common operation on calculating the expectation of some function ϕ with respect to a given distribution. Let $p(\Theta) = \bar{p}(\Theta)/Z$ be the density of a probability distribution that we compute $\bar{p}(\Theta)$ pointwise up to a normalizing constant Z . The expectation of interest is

$$I = \int \phi(\Theta)p(\Theta)d\Theta. \quad (8)$$

Replacing $p(\cdot)$ by $\hat{p}(\cdot)$, we get the unbiased Monte Carlo estimate of this quantity:

$$\hat{I}_{\text{MC}} = \frac{1}{N} \sum_{i=1}^N \phi(\Theta_i). \quad (9)$$

Asymptotically, when $N \rightarrow \infty$ the estimate \hat{I}_{MC} will almost surely converge to I by the strong law of large numbers. In practice, however, we often cannot sample from p directly. Many methods have been developed, such as rejection sampling and importance sampling, which however often suffer from severe limitations in high dimensional spaces. We refer the readers to the book [156] and the review article [15] for details. Below, we introduce Markov chain Monte Carlo (MCMC), a very general and powerful framework that allows sampling from a broad family of distributions and scales well with the dimensionality of the sample space. More importantly, many advances have been made on scalable MCMC methods for Big Data, which will be discussed later.

An MCMC method constructs an ergodic p -stationary Markov chain sequentially. Once the chain has converged (i.e., finishing the burn-in phase), we can use the samples to estimate I . The Metropolis-Hastings algorithm [125], [83] constructs such a chain by using the following rule to transit from the current state Θ_t to the next state Θ_{t+1} :

- 1) draw a candidate state Θ' from a proposal distribution $q(\Theta|\Theta_t)$;
- 2) compute the acceptance probability:

$$A(\Theta', \Theta_t) \triangleq \min \left(1, \frac{\bar{p}(\Theta')q(\Theta_t|\Theta')}{\bar{p}(\Theta_t)q(\Theta'|\Theta_t)} \right). \quad (10)$$

- 3) draw $\gamma \sim \text{Uniform}[0,1]$. If $\gamma < A(\Theta', \Theta_t)$ set $\Theta_{t+1} \leftarrow \Theta'$, otherwise set $\Theta_{t+1} \leftarrow \Theta_t$.

Note that for Bayesian models, each MCMC step involves an evaluation of the full likelihood to get the (unnormalized) posterior $\bar{p}(\Theta)$, which can be prohibitive for big learning with massive data sets. We will revisit this problem later.

One special type of MCMC methods is the Gibbs sampling [69], which iteratively draws samples from local conditionals. Let Θ be a M -dimensional vector. The standard Gibbs sampler performs the following steps to get a new sample $\Theta^{(t+1)}$:

- 1) draw a sample $\theta_1^{(t+1)} \sim p(\theta_1|\theta_2^{(t)}, \dots, \theta_M^{(t)})$;
- 2) for $j = 2 : M - 1$, draw a sample

$$\theta_j^{(t+1)} \sim p(\theta_j|\theta_1^{(t+1)}, \dots, \theta_{j-1}^{(t+1)}, \theta_{j+1}^{(t)}, \dots, \theta_M^{(t)});$$

- 3) draw a sample $\theta_M^{(t+1)} \sim p(\theta_M|\theta_1^{(t+1)}, \dots, \theta_{M-1}^{(t+1)})$;

One issue with MCMC methods is that the convergence rate can be prohibitively slow even for conventional applications. Extensive efforts have been spent to improve the convergence rates. For example, hybrid Monte Carlo methods explore gradient information to improve the mixing rates when the model parameters are continuous, with representative examples of Langevin dynamics and Hamiltonian dynamics [136]. Other improvements include population-based MCMC methods [89] and annealing methods [73] that can sometimes handle distributions

with multiple modes. Another useful technique to develop simpler or more efficient MCMC methods is data augmentation [174], [62], [135], which introduces auxiliary variables to transform marginal dependency into a set of conditional independencies. For Gibbs samplers, the useful techniques to improve the convergence include blockwise Gibbs sampling and partially collapsed Gibbs (PCG) sampling [181]. A PCG sampler is as simple as an ordinary Gibbs sampler, but often improves the convergence by replacing some of the conditional distributions of an ordinary Gibbs sampler with conditional distributions of some marginal distributions.

2.4 FAQ

Common questions regarding Bayesian methods are:

Q: Why should I use Bayesian methods?

A: There are many reasons for choosing Bayesian methods, as discussed in the Introduction. A formal theoretical argument is provided by the classic de Finetti theorem, which states that: If $(\mathbf{x}_1, \mathbf{x}_2, \dots)$ are infinitely exchangeable, then for any N

$$p(\mathbf{x}_1, \dots, \mathbf{x}_N) = \int \left(\prod_{i=1}^N p(\mathbf{x}_i|\theta) \right) dP(\theta) \quad (11)$$

for some random variable θ and probability measure P . The infinite exchangeability is an often satisfied property. For example, any i.i.d data are infinitely exchangeable. Moreover, the data whose ordering information is not informative is also infinitely exchangeable, e.g., the commonly used bag-of-words representation of documents [36] and images [112].

Q: How should I choose the prior?

A: There are two schools of thought, namely, objective Bayes and subjective Bayes. For objective Bayes, an improper noninformative prior (e.g., the Jeffreys prior [91] and the maximum-entropy prior [90]) is used to capture ignorance, which admits good frequentist properties. In contrast, subjective Bayesian methods embrace the influence of priors. A prior may have some parameters λ . Since it is often difficult to elicit an honest prior, e.g., setting the true value of λ , two practical methods are often used. One is hierarchical Bayesian methods, which assume a hyper-prior on λ and define the prior as a marginal distribution:

$$p_0(\Theta) = \int p_0(\Theta|\lambda)p(\lambda)d\lambda. \quad (12)$$

Though $p(\lambda)$ may have hyper-parameters as well, it is commonly believed that these parameters will have a weak influence as long as they are far from the likelihood model, thus can be fixed at some convenient values or put another layer of hyper-prior.

Another method is *empirical Bayes*, which adopts a data-driven estimate $\hat{\lambda}$ and uses $p_0(\Theta|\hat{\lambda})$ as the prior. Empirical Bayes can be seen as an approximation

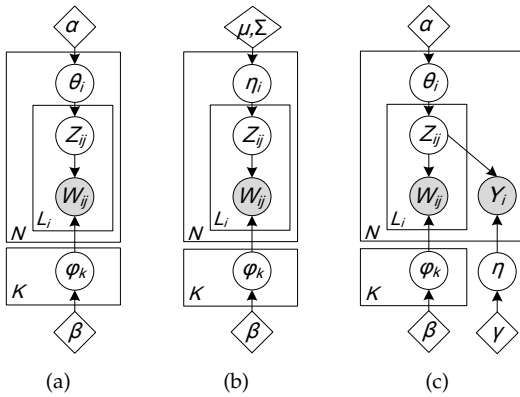


Fig. 1. Graphical illustration of (a) LDA [36]; (b) logistic-normal topic model [34]; and (c) supervised LDA.

to the hierarchical approach, where $p(\lambda)$ is approximated by a delta-Dirac mass $\delta_{\hat{\lambda}}(\lambda)$. One common choice is maximum marginal likelihood estimate, that is, $\hat{\lambda} = \operatorname{argmax}_{\lambda} p(\mathcal{D}|\lambda)$. Empirical Bayes has been applied in many problems, including variable selection [70] and nonparametric Bayesian methods [123]. Recent progress has been made on characterizing the conditions when empirical Bayes merges with the Bayesian inference [148] as well as the convergence rates of empirical Bayes methods [58].

In practice, another important consideration is the tradeoff between model capacity and computational cost. If a prior is conjugate to the likelihood, the posterior inference will be relatively simpler in terms of computation and memory demands, as the posterior belongs to the same family as the prior.

Example 1: Dirichlet-Multinomial Conjugate Pair

Let $\mathbf{x} \in \{0,1\}^V$ be a one-hot representation of a discrete variable with V possible values. It is easy to verify that for the multinomial likelihood, $p(\mathbf{x}|\boldsymbol{\theta}) = \prod_{k=1}^V \theta_k^{x_k}$, the conjugate prior is a Dirichlet distribution, $p_0(\boldsymbol{\theta}|\boldsymbol{\alpha}) = \operatorname{Dir}(\boldsymbol{\alpha}) = \frac{1}{Z} \prod_{k=1}^V \theta_k^{\alpha_k-1}$, where $\boldsymbol{\alpha}$ is the hyper-parameter and Z is the normalization factor. In fact, the posterior distribution is $\operatorname{Dir}(\boldsymbol{\alpha} + \mathbf{x})$.

A popular Bayesian model that explores such conjugacy is latent Dirichlet allocation (LDA) [36], as illustrated in Fig. 1(a). LDA posits that each document \mathbf{w}_i is an admixture of a set of K topics, of which each topic ψ_k is a unigram distribution over a given vocabulary. The generative process is as follows:

- 1) draw K topics $\psi_k \sim \operatorname{Dir}(\boldsymbol{\beta})$
- 2) for each document $i \in [N]$:
 - a) draw a topic mixing vector $\boldsymbol{\theta}_i \sim \operatorname{Dir}(\boldsymbol{\alpha})$
 - b) for each word $j \in [L_i]$ in document i :
 - i) draw a topic assignment $z_{ij} \sim \operatorname{Multi}(\boldsymbol{\theta}_i)$
 - ii) draw a word $w_{ij} \sim \operatorname{Multi}(\psi_{z_{ij}})$.

LDA has been popular in many applications. However, a conjugate prior can be restrictive. For example, the Dirichlet distribution does not impose correlation between different parameters, except the normalization constraint. In order to obtain more flexible models, a non-conjugate prior can be chosen.

Example 2: Logistic-Normal Prior A logistic-normal distribution [11] provides one way to impose correlation structure among the multiple dimensions of $\boldsymbol{\theta}$. It is defined as follows:

$$\boldsymbol{\eta} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \quad \theta_k = \frac{e^{\eta_k}}{\sum_j e^{\eta_j}}. \quad (13)$$

This prior has been used to develop correlated topic models (or logistic-normal topic models) [34], which can infer the correlation structure among topics. However, the flexibility pays cost on computation, needing scalable algorithms to learn large topic graphs [48].

3 BIG BAYESIAN LEARNING

Though much more emphasis in big Bayesian learning has been put on scalable algorithms and systems, substantial advances have been made on *adaptive* and *flexible* Bayesian methods. This section reviews nonparametric Bayesian methods for adaptively inferring model complexity and regularized Bayesian inference for improving the flexibility via posterior regularization, while leaving the large part of scalable algorithms and systems to next sections.

3.1 Nonparametric Bayesian Methods

For parametric Bayesian models, the parameter space is pre-specified. No matter how the data changes, the number of parameters is fixed. This restriction may cause limitations on model capacity, especially for big data applications, where it may be difficult or even counter-productive to fix the number of parameters a priori. For example, a Gaussian mixture model with a fixed number of clusters may fit the given data set well; however, it may be sub-optimal to use the same number of clusters if more data comes under a slightly changed distribution. It would be ideal if the clustering models can figure out the unknown number of clusters automatically. Similar requirements on automatic model selection exist in feature representation learning [29] or factor analysis, where we would like the models to automatically figure out the dimension of latent features (or factors) and maybe also the topological structure among features (or factors) at different abstraction levels [7].

Nonparametric Bayesian (NPB) methods provide an elegant solution to such needs on automatic adaptation of model capacity when learning a single model. Such adaptivity is obtained by defining stochastic processes on rich measure spaces. Classical examples include Dirichlet process (DP), Indian buffet process (IBP), and Gaussian process (GP). Below, we briefly review DP and IBP. We refer the readers to the articles [74], [72], [132] for a nice overview and the textbook [85] for a comprehensive treatment.

3.1.1 Dirichlet Process

A DP defines the distribution of random measures. It was first developed in [65]. Specifically, a DP is parameterized by a concentration parameter $\alpha > 0$ and a base distribution G_0 over a measure space Ω . A random variable drawn from a DP, $G \sim \mathcal{DP}(\alpha, G_0)$, is itself a distribution over Ω . It was shown that the random distributions drawn from a DP are discrete almost surely, that is, they place the probability mass on a countably infinite collection of atoms, i.e.,

$$G = \sum_{k=1}^{\infty} \pi_k \delta_{\theta_k}, \quad (14)$$

where θ_k is the value (or location) of the k th atom independently drawn from the base distribution G_0 and π_k is the probability assigned to the k th atom. Sethuraman [166] provided a constructive definition of π_k based on a stick-breaking process as in Fig. 2(a). Consider a stick with unit length. We break the stick into an infinite number of segments π_k by the following process with $\nu_k \sim \text{Beta}(1, \alpha)$:

$$\pi_1 = \nu_1, \quad \pi_k = \nu_k \prod_{j=1}^{k-1} (1 - \nu_j), \quad k = 2, 3, \dots, \infty. \quad (15)$$

That is, we first choose a beta variable ν_1 and break ν_1 of the stick. Then, for the remaining segment, we draw another beta variable and break off that proportion of the remainder of the stick. Such a representation of DP provides insights for developing variational approximate inference algorithms [33].

DP is closely related to the Chinese restaurant process (CRP) [150], which defines a distribution over infinite partitions of integers. CRP derives its name from a metaphor: Image a restaurant with an infinite number of tables and a sequence of customers entering the restaurant and sitting down. The first customer sits at the first table. For each of the subsequent customers, she sits at each of the occupied tables with a probability proportional to the number of previous customers sitting there, and at the next unoccupied table with a probability proportional to α . In this process, the assignment of customers to tables defines a random partition. In fact, if we repeatedly draw a set of samples from G , that is, $\theta_i \sim G$, $i \in [N]$, then it was shown that the joint distribution of $\theta_{1:N}$

$$p(\theta_1, \dots, \theta_N | \alpha, G_0) = \int \left(\prod_{i=1}^N p(\theta_i | G) \right) dP(G | \alpha, G_0)$$

exists a clustering property, that is, the θ_i s will share repeated values with a non-zero probability. These shared values define a partition of the integers from 1 to N , and the distribution of this partition is a CRP with parameter α . Therefore, DP is the de Finetti mixing distribution of CRP.

Antoniak [17] first developed DP mixture models by adding a data generating step, that is, $\mathbf{x}_i \sim$




Fig. 2. The stick-breaking process for: (a) DP; (b) IBP.

$p(\mathbf{x} | \theta_i)$, $i \in [N]$. Again, marginalizing out the random distribution G , the DP mixture reduces to a CRP mixture, which enjoys nice Gibbs sampling algorithms [134]. For DP mixtures, a slice sampler [135] has been developed [184], which transforms the infinite sum in Eq. (14) into a finite sum conditioned on some uniformly distributed auxiliary variable.

3.1.2 Indian Buffet Process

A mixture model assumes that each data is assigned to one single component. Latent factor models weaken this assumption by associating each data with some or all of the components. When the number of components is smaller than the feature dimension, latent factor models provide dimensionality reduction. Popular examples include factor analysis, principal component analysis and independent component analysis. The general assumption of a latent factor model is that the observed data $\mathbf{x} \in \mathbb{R}^P$ is generated by a noisy weighted combination of latent factors, that is,

$$\mathbf{x}_i = W \mathbf{z}_i + \epsilon_i, \quad (16)$$

where W is a $P \times K$ factor loading matrix, with element W_{mk} expressing how latent factor k influences the observation dimension m ; \mathbf{z}_i is a K -dimensional vector expressing the activity of each factor; and ϵ_i is a vector of independent noise terms (usually Gaussian noise). In the above models, the number of factors K is assumed to be known. Indian buffet process (IBP) [80] provides a nonparametric Bayesian variant of latent factor models and it allows the number of factors to grow as more data are observed.

Consider binary factors for simplicity². Putting the latent factors of N data points in a matrix Z , of which the i th row is \mathbf{z}_i . IBP defines a process over the space of binary matrixes with an unbounded number of columns. IBP derives its name from a similar metaphor as CRP. Image a buffet with an infinite number of dishes (factors) arranged in a line and a sequence of customers choosing the dishes. Let z_{ik} denote whether customer i chooses dish k . Then, the first customer chooses K_1 dishes, where $K_1 \sim \text{Poisson}(\alpha)$; and the subsequent customer n (> 1) chooses:

- 1) each of the previously sampled dishes with probability m_k/n , where m_k is the number of customers who have chosen dish k ;

2. Real-valued factors can be easily considered by defining $\mathbf{h}_i = \mathbf{z}_i \odot \boldsymbol{\mu}_i$, where the binary \mathbf{z}_i are 0/1 masks to indicate whether a factor is active or not, and $\boldsymbol{\mu}_i$ are the values of the factors.

2) K_i additional dishes, where $K_i \sim \text{Poisson}(\alpha/n)$.

IBP plays the same role for latent factor models that CRP plays for mixture models, allowing an unbounded number of latent factors. Analogous to the role that DP is the de Finetti mixing distribution of CRP, the de Finetti mixing distribution underlying IBP is a Beta process [178]. IBP also admits a stick-breaking representation [175] as shown in Fig. 2(b), where the stick lengths are defined as

$$\nu_k \sim \text{Beta}(\alpha, 1), \quad \pi_k = \prod_{j=1}^k \nu_j, \quad k = 1, 2, \dots, \infty. \quad (17)$$

Note that unlike the stick-breaking representation of DP, where the stick lengths sum to 1, the stick lengths here need not sum to 1. Such a representation has led to the developments of Monte Carlo [175] as well as variational approximation inference algorithms [59].

3.1.3 Extensions

To meet the flexibility and adaptivity requirements of big learning, many recent advances have been made on developing sophisticated NPB methods for modeling various types of data, including grouped data, spatial data, time series, and networks.

Hierarchical models are natural tools to describe grouped data, e.g., documents from different source domains. Hierarchical Dirichlet process (HDP) [176] and hierarchical Beta process [178] have been developed, allowing an infinite number of latent components to be shared by multiple domains. The work [7] presents a cascading IBP (CIBP) to learn the topological structure of multiple layers of latent features, including the number of layers, the number of hidden units at each layer, the connection structure between units at neighboring layers, and the activation function of hidden units. The recent work [53] presents an extended CIBP process to generate connections between non-consecutive layers.

Another dimension of the extensions concerns modeling the dependencies between observations in a time series. For example, DP has been used to develop the infinite hidden Markov models [23], which posit the same sequential structure as in the hidden Markov models, but allowing an infinite number of latent classes. In [176], it was shown that iHMM is a special case of HDP. The recent work [203] presents a max-margin training of iHMMs under the regularized Bayesian framework, as will be reviewed shortly.

Finally, for spatial data, modeling dependency between nearby data points is important. Recent extensions of Bayesian nonparametric methods include the dependent Dirichlet process [118], spatial Dirichlet process [60], distance dependent CRP [32], dependent IBP [194], and distance dependent IBP [71]. For network data analysis (e.g., social networks, biological networks, and citation networks), recent extensions include the nonparametric Bayesian relational latent

feature models for link prediction [126], [205], which adopt IBP to allow for an unbounded number of latent features, and the nonparametric mixed membership stochastic block models for community discovery [78], [99], which use HDP to allow mixed membership in an unbounded number of latent communities.

3.2 Regularized Bayesian Inference

Regularized Bayesian inference (RegBayes) [209] represents one recent advance that extends the scope of Bayesian methods on incorporating rich side information. Recall that the classic Bayes' theorem is equivalent to a variational optimization problem as in (2). RegBayes builds on this formulation and defines the generic optimization problem

$$\min_{q(\Theta) \in \mathcal{P}} \text{KL}(q(\Theta) \| p(\Theta | \mathcal{D})) + c \cdot \Omega(q(\Theta); \mathcal{D}), \quad (18)$$

where $\Omega(q(\Theta); \mathcal{D})$ is the *posterior regularization* term; c is a nonnegative regularization parameter; and $p(\Theta | \mathcal{D})$ is the ordinary Bayesian posterior. Fig. 3 provides a high-level comparison between RegBayes and Bayes' rule. Several questions need to be answered in order to solve practical problems.

Q: How to define the posterior regularization?

A: In general, posterior regularization can be any informative constraints that are expected to regularize the properties of the posterior distribution. It can be defined as the large-margin constraints to enforce a good prediction accuracy [206], or the logic constraints to incorporate expert knowledge [124], or the sparsity constraints [103].

Example 3: Max-margin LDA Following the paradigm of ordinary Bayes, a supervised topic model is often defined by augmenting the likelihood model. For example, the supervised LDA (sLDA) [35] has a similar structure as LDA (see Fig. 1(c)), but with an additional likelihood $p(y_d | \mathbf{z}_d, \boldsymbol{\eta})$ to describe labels. Such a design can lead to an imbalanced combination of the word likelihood $p(\mathbf{w}_d | \mathbf{z}_d, \boldsymbol{\psi})$ and the label likelihood because a document often has tens or hundreds of words while only one label. The imbalance problem causes unsatisfactory prediction results [210].

To improve the discriminative power of supervised topic models, the max-margin MedLDA has been developed, under the RegBayes framework. Consider binary classification for simplicity. In this case, we have $\Theta = \{\boldsymbol{\theta}_i, \mathbf{z}_i, \boldsymbol{\psi}_k\}$. Let $f(\boldsymbol{\eta}, \mathbf{z}_i) = \boldsymbol{\eta}^\top \bar{\mathbf{z}}_i$ be the discriminant function³, where $\bar{\mathbf{z}}_i$ is the average topic assignments, with $\bar{z}_i^k = \frac{1}{L_i} \sum_j \mathbb{I}(z_{ij} = k)$. The posterior regularization can be defined in two ways:

Averaging classifier: An averaging classifier makes predictions using the expected discriminant function, that is, $\hat{y}(q) = \text{sign}(\mathbb{E}_q[f(\boldsymbol{\eta}, \mathbf{z})])$. Then, the posterior

3. We ignore the offset for simplicity.




Fig. 3. (a) Bayesian inference with the Bayes' rule; and (b) regularized Bayesian inference (RegBayes) which solves an optimization problem with a posterior regularization term to incorporate rich side information.

regularization

$$\Omega^{\text{Avg}}(q(\Theta); \mathcal{D}) = \sum_{i=1}^N (1 - y_i \mathbb{E}_q[f(\eta, \mathbf{z}_i)])_+$$

is an upper bound of the training error, therefore a good surrogate loss for learning, where $(x)_+ = \max(0, x)$. This strategy has been adopted in MedLDA [206].

Gibbs classifier: A Gibbs classifier (or stochastic classifier) randomly draws a sample (η, \mathbf{z}_d) from the target posterior $q(\Theta)$ and makes predictions using the latent prediction rule, that is, $\hat{y}(\eta, \mathbf{z}_i) = \text{sign}f(\eta, \mathbf{z}_i)$. Then, the posterior regularization is defined as

$$\Omega^{\text{Gibbs}}(q(\Theta); \mathcal{D}) = \mathbb{E}_q \left[\sum_{i=1}^N (1 - y_i f(\eta, \mathbf{z}_i))_+ \right].$$

This strategy has been adopted to develop Gibbs MedLDA [207].

The two strategies are closely related, e.g., we can show that $\Omega^{\text{Gibbs}}(q(\Theta))$ is an upper bound of $\Omega^{\text{Avg}}(q(\Theta))$. The formulation with a Gibbs classifier can lead to a scalable Gibbs sampler by using data augmentation techniques [211]. If a logistic log-loss is adopted to define the posterior regularization, an improved sLDA model can be developed to address the imbalance issue and lead to significantly more accurate predictions [210].

Q: What is the relationship between prior, likelihood, and posterior regularization?

A: Though the three part are closely connected, there are some key differences. First, prior is chosen before seeing data, while both likelihood and posterior regularization depend on the data. Second, different from the likelihood, which is restricted to be a normalized distribution, no constraints are imposed on the posterior regularization. Therefore, posterior regularization is much more flexible than prior or likelihood. In fact, it can be shown that (1) putting constraints on priors is a special case of posterior regularization, where the regularization term does not depend on data; and (2) RegBayes can be more flexible than standard Bayes' rule, that is, there exists some RegBayes posterior distributions that are not achievable by the Bayes' rule [209].

Q: How to solve the optimization problem?

A: The posterior regularization term affects the difficulty of solving problem (18). When the regularization term is a convex functional of $q(\Theta)$, which is common in many applications such as the above max-margin formulations, the optimal solution can be characterized in a general form via convex duality theory [209]. When the regularization term is non-convex, a generalized representation theorem can also be derived, but requires more effects on dealing with the non-convexity [103].

4 SCALABLE ALGORITHMS

To deal with big data, the posterior inference algorithms should be scalable. Significant advances have been made in two aspects: (1) using random sampling to do stochastic or online Bayesian inference; and (2) using multi-core and multi-machine architectures to do parallel and distributed Bayesian inference.

4.1 Stochastic Algorithms

In Big Learning, the intriguing results of [38] suggest that an algorithm as simple as stochastic gradient descent (SGD) can be optimally efficient in terms of “number of bits learned per unit of computation”. For Bayesian models, both stochastic variational and stochastic Monte Carlo methods have been developed to explore the redundancy of data relative to a model by subsampling data examples for every update and reasoning about the uncertainty created in this process [190]. We overview each type in turn.

4.1.1 Stochastic Variational Methods

As we have stated in Section 2.3.1, variational methods solve an optimization problem to find the best approximate distribution to the target posterior. When the variational distribution is characterized in some parametric form, this problem can be solved with stochastic gradient descent (SGD) methods [40] or the adaptive SGD [61]. A SGD method randomly draws a subset B_t and updates the variational parameters using the estimated gradients, that is,

$$\phi_{t+1} \leftarrow \phi_t + \epsilon_t (\nabla_{\phi} \text{KL}(q||p_0(\theta)) - \nabla_{\phi} \mathbb{E}_q[\log p(\mathcal{D}|\theta)]),$$

where the full data gradient is approximated as

$$\nabla_{\phi} \mathbb{E}_q[\log p(\mathcal{D}|\theta)] \approx \frac{N}{|B_t|} \sum_{i \in B_t} \nabla_{\phi} \mathbb{E}_q[\log p(\mathbf{x}_i|\theta)], \quad (19)$$

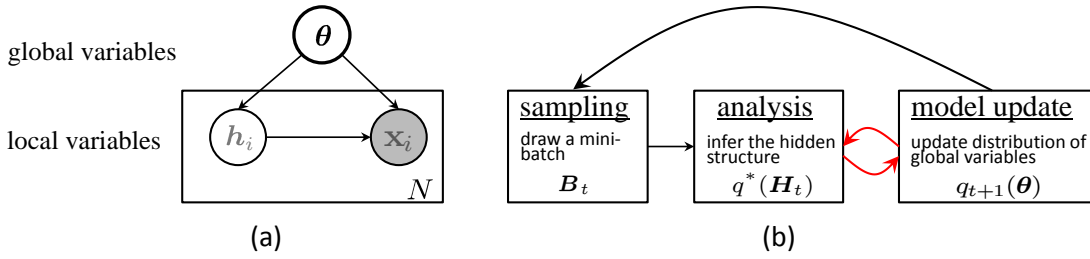


Fig. 4. (a) The general structure of Bayesian latent variable models, where \mathbf{h}_d denotes the local latent variables for each data d ; (b) the process of stochastic variational inference, where the red arrows denote that in practice we may need multiple iterations between “analysis” and “model update” to have fast convergence.

and ϵ_t is a learning rate. If the noisy gradient is an unbiased estimate of the true gradient, the procedure is guaranteed to approach the optimal solution when the learning rate is appropriately set [37].

For Bayesian latent variable models, we need to infer the latent variables when performing the updates. In general, we can group the latent variables into two categories — global variables and local variables. Global variables correspond to the model parameters θ (e.g., the topics ψ in LDA), while local variables represent some hidden structures of the data (e.g., the topic assignments \mathbf{z} in an LDA with the topic mixing proportions collapsed out). Fig. 4 provides an illustration of such models and the stochastic variational inference, which consists of three steps:

- 1) randomly draw a mini-batch B_t of data samples;
- 2) infer the local latent variables for each data in B_t ;
- 3) update the global variables.

However, the standard gradients over the parameters ϕ may not be the most informative direction (i.e., the steepest direction) to search for the distribution q . A better way is to use natural gradient [13], which is the steepest search direction in a Riemannian manifold space of probability distributions [87]. To reduce the efforts on hand-tuning the learning rate, which often influences the performance much, the work [153] presents an adaptive learning rate while [169] adopts Bayesian optimization to search for good learning rates, both learning to faster convergence. By borrowing the gradient averaging ideas from stochastic optimization, [121] proposes to use smoothed gradients in stochastic variational inference to reduce the variance (by trading-off the bias). Stochastic variational inference methods have been studied for many Bayesian models, such as topic models and hierarchical Dirichlet process [87].

In many cases, the ELBO and its gradient may be intractable to compute due to the intractability of the expectation over variational distributions. Two types of methods are commonly used to address this problem. First, another layer of variational bound is derived by introducing additional variational parameters. This has been used in many examples, such as the logistic-normal topic models [34] and supervised LDA [35]. For such methods, it is important to develop tight variational bounds for specific models [122], which

is still an active area. Another type of methods is to use Monte Carlo estimates of the variational bound as well as its gradients. Recent work includes the stochastic approximation scheme with variance reduction [143], [153] and the auto-encoding variational Bayes (AEVB) [100] that learns a neural network (a.k.a recognition model) to represent the variational distribution for continuous latent variables.

Consider the model with one layer of continuous latent variables \mathbf{h}_i in Fig. 4 (a). Assume the variational distribution $q_\phi(\Theta) = q_\phi(\theta) \prod_{i=1}^N q_\phi(\mathbf{h}_i|\mathbf{x}_i)$. Let $G_\phi(\mathbf{x}, \mathbf{h}, \theta) = \log p(\mathbf{h}|\theta) + \log p(\mathbf{x}|\mathbf{h}, \theta) - \log q_\phi(\mathbf{h}|\mathbf{x})$. The ELBO in Eq. (2) can be written as

$$\mathcal{L}(\phi; \mathcal{D}) = \mathbb{E}_q \left[\log p_0(\theta) + \sum_i G_\phi(\mathbf{x}_i, \mathbf{h}_i, \theta) - \log q_\phi(\theta) \right].$$

By using the equality $\nabla_\phi q_\phi(\Theta) = q_\phi(\Theta) \nabla_\phi \log q_\phi(\Theta)$, it can be shown that the gradient is

$$\nabla_\phi \mathcal{L} = \mathbb{E}_q [(\log p(\Theta, \mathcal{D}) - \log q_\phi(\Theta)) \nabla_\phi \log q_\phi(\Theta)].$$

A naive Monte Carlo estimate of the gradient is

$$\nabla_\phi \mathcal{L} \approx \frac{1}{L} \sum_{l=1}^L [(\log p(\Theta^l, \mathcal{D}) - \log q_\phi(\Theta^l)) \nabla_\phi \log q_\phi(\Theta^l)],$$

where $\Theta^l \sim q_\phi(\Theta)$. Note that the sampling and the gradient $\nabla_\phi \log q_\phi(\Theta^l)$ only depend on the variational distribution, not the underlying model. However, the variance of such an estimate can be too large to be useful. In practice, effective variance reduction techniques are needed [143], [153].

For continuous \mathbf{h} , a reparameterization of the samples $\mathbf{h} \sim q_\phi(\mathbf{h}|\mathbf{x})$ can be derived using a differentiable transformation $g_\phi(\epsilon, \mathbf{x})$ of a noise variable ϵ :

$$\mathbf{h} = g_\phi(\epsilon, \mathbf{x}), \text{ where } \epsilon \sim p(\epsilon). \quad (20)$$

This is known as *non-centered parameterization* (NCP) in statistics [145], while the original representation is known as *centered parameterization* (CP). A similar NCP reparameterization exists for the continuous θ :

$$\theta = f_\phi(\zeta), \text{ where } \zeta \sim p(\zeta). \quad (21)$$

Given a minibatch of data points B_t , we define $F_\phi(\{\mathbf{x}_i, \mathbf{h}_i\}_{i \in B_t}, \theta) = \frac{N}{|B_t|} \sum_{i \in B_t} G_\phi(\mathbf{x}_i, \mathbf{h}_i, \theta) +$

$\log p_0(\theta) - \log q_\phi(\theta)$. Then, the Monte Carlo estimate of the variational lower bound is

$$\mathcal{L}(\phi; \mathcal{D}) \approx \frac{1}{L} \sum_{l=1}^L F_\phi \left(\{\mathbf{x}_i, g_\phi(\epsilon^l, \mathbf{x}_i)\}_{i \in B_l}, f_\phi(\zeta^l) \right), \quad (22)$$

where $\epsilon^l \sim p(\epsilon)$ and $\zeta^l \sim p(\zeta)$. This stochastic estimate can be maximized via gradient ascent methods.

It has been analyzed that CP and NCP possess complimentary strengths [145], in the sense that NCP is likely to work when CP does not and conversely. An accompany paper [101] to AEVB analyzes the conditions for gradient-based samplers (e.g., HMC) when such a differentiable NCP can be effective or ineffective in reducing posterior dependencies, and suggests to use the interleaving strategy between centered and non-centered parameterization as previously studied in [200]. AEVB has been extended to learn deep generative models [155] using the similar reparameterization trick on continuous latent variables, while [130] presents a sophisticated method to reduce the variance of the naive Monte Carlo estimate for deep autoregressive models, which is applicable to both continuous and discrete latent variables.

4.1.2 Stochastic Monte Carlo Methods

The existing stochastic Monte Carlo methods can be generally grouped into three categories, namely, stochastic gradient-based methods, the methods using approximate MH test with randomly sampled mini-batches, and data augmentation.

Stochastic Gradient: The idea of using gradient information to improve the mixing rates has been systematically studied in various MC methods, including Langevin dynamics and Hamiltonian dynamics [136]. For example, the Langevin dynamics is an MCMC method that produces samples from the posterior by means of gradient updates plus Gaussian noise, resulting in a proposal distribution $p(\theta_{t+1}|\theta_t)$ by the following equation:

$$\theta_{t+1} = \theta_t + \frac{\epsilon_t}{2} (\nabla_\theta \log p_0(\theta) + \nabla_\theta \log p(\mathcal{D}|\theta)) + \zeta_t, \quad (23)$$

where $\zeta_t \sim \mathcal{N}(0, \epsilon_t I)$ is an isotropic Gaussian noise and $\log p(\mathcal{D}|\theta) = \sum_i \log p(\mathbf{x}_i|\theta)$ is the log-likelihood of the full data set. The mean of the proposal distribution is in the direction of increasing log posterior due to the gradient, while the Gaussian noise will prevent the samples from collapsing to a single maximum. A Metropolis-Hastings correction step is required to correct for discretisation error [158].

The stochastic ideas have been successfully explored in these methods to develop efficient stochastic Monte Carlo methods, including stochastic gradient Langevin dynamics (SGLD) [191] and stochastic gradient Hamiltonian dynamics (SGHD) [49]. For example, SGLD replaces the calculation of the gradient over the full data set, with a stochastic approximation based on a subset of data. Let B_t be the subset of

data points uniformly sampled from the full data set at iteration t . Then, the gradient is approximated as:

$$\nabla_\theta \log p(\mathcal{D}|\theta) \approx \frac{N}{|B_t|} \sum_{i \in B_t} \nabla_\theta \log p(\mathbf{x}_i|\theta). \quad (24)$$

Note that SGLD doesn't use a MH correction step, as calculating the acceptance probability requires use of the full data set. Convergence to the posterior is still guaranteed if the step sizes are annealed to zero at a certain rate, as rigorously justified in [149], [177].

To further improve the mixing rates, the stochastic gradient Fisher scoring method [9] was developed, which represents an extension of the Fisher scoring method based on stochastic gradients [163] by incorporating randomness in a subsampling process. Similarly, exploring the Riemannian manifold structure leads to the development of stochastic gradient Riemannian Langevin dynamics (SGRLD) [147], which performs SGLD on the probability simplex space.

Approximate MH Test: Another category of stochastic Monte Carlo methods rely on approximate MH test using randomly sampled subset of data points, since an exact calculation of the MH test in Eq. (10) scales linearly to the data size, which is prohibitive for large-scale data sets. For example, the work [102] presents an approximate MH rule via sequential hypothesis testing, which allows us to accept or reject samples with high confidence using only a fraction of the data required for the exact MH rule. The systematic bias and its tradeoff with variance were theoretically analyzed. Specifically, it is based on the observation that the MH test rule in Eq. (10) can be equivalently written as follows:

- 1) Draw $\gamma \sim \text{Uniform}[0, 1]$ and compute:

$$\begin{aligned} \mu_0 &= \frac{1}{N} \log \left[\gamma \frac{p_0(\Theta_t)q(\Theta'| \Theta_t)}{p_0(\Theta')q(\Theta_t | \Theta')} \right] \\ \mu &= \frac{1}{N} \sum_{i=1}^N \ell_i, \text{ where } \ell_i = \log \frac{p(\mathbf{x}_i|\Theta')}{p(\mathbf{x}_i|\Theta_t)}; \end{aligned}$$

- 2) If $\mu > \mu_0$ set $\Theta_{t+1} \leftarrow \Theta'$; otherwise $\Theta_{t+1} \leftarrow \Theta_t$.

Note that μ_0 is independent of the data set, thus can be easily calculated. This reformulation of the MH test makes it very easy to frame it as a statistical hypothesis test, that is, given μ_0 and a set of samples $\{\ell_{t_1}, \dots, \ell_{t_n}\}$ drawn without replacement from the population $\{\ell_1, \dots, \ell_N\}$, can we decide whether the population mean μ is greater than or less than the threshold μ_0 ? Such a test can be done by increasing the cardinality of the subset until a prescribed confidence level is reached. The MH test with approximate confidence intervals can be combined with the above stochastic gradient methods (e.g., SGLD) to correct their bias. The similar sequential testing ideas can be applied to Gibbs sampling, as discussed in [102].

Under the similar setting of approximate MH test with subsets of data, the work [21] derives a new

stopping rule based on some concentration bounds (e.g., the empirical Bernstein bound [22]), which leads to an adaptive sampling strategy with theoretical guarantees on the total variational norm between the approximate MH kernel and the target distribution of MH applied to the full data set.

Data Augmentation: The work [119] presents a Firefly Monte Carlo (FlyMC) method, which is guaranteed to converge to the true target posterior. FlyMC relies on a novel data augmentation formulation [62]. Specifically, let z_i be a binary variable, indicating whether data i is active or not, and $B_i(\Theta)$ be a strictly positive lower bound of the i th likelihood: $0 < B_i(\Theta) < L_i(\Theta) \triangleq p(\mathbf{x}_i|\Theta)$. Then, the target posterior $p(\Theta|\mathcal{D})$ is the marginal of the complete posterior with the augmented variables $\mathbf{Z} = \{z_i\}_{i=1}^N$:

$$p(\Theta, \mathbf{Z}|\mathcal{D}) \propto p_0(\Theta) \prod_{i=1}^N p(\mathbf{x}_i|\Theta) p(z_i|\mathbf{x}_i, \Theta), \quad (25)$$

where $p(z_i|\mathbf{x}_i, \Theta) = (1 - \gamma_i)^{z_i} \gamma_i^{(1-z_i)}$ and $\gamma_i = B_i(\Theta)/L_i(\Theta)$. Then, we can construct a Markov chain for the complete posterior by alternating between updates of Θ conditioned on \mathbf{Z} , which can be done with any conventional MCMC algorithm, and updates of \mathbf{Z} conditioned on Θ , which can also be efficiently done as we only need to re-calculate the likelihoods of the data points with active z variables, thus effectively using a random subset of data points in each iteration of the MC methods.

4.2 Streaming Algorithms

We can see that both (19) and (24) need to know the data size N , which renders them unsuitable for learning with streaming data, where data comes in small batches without an explicit bound on the total number as times goes along, e.g., tracking an aircraft using radar measurements. This conflicts with the sequential nature of the Bayesian updating procedure. Specifically, let B_t be the small batch at time t . Given the posterior at time t , $p_t(\Theta) := p(\Theta|B_1, \dots, B_t)$, the posterior distribution at time $t+1$ is

$$p_{t+1}(\Theta) := p(\Theta|B_1, \dots, B_{t+1}) = \frac{p_t(\Theta)p(B_{t+1}|\Theta)}{p(B_1, \dots, B_{t+1})}. \quad (26)$$

In other words, the posterior at time t is actually playing the role of a prior for the data at time $t+1$ for the Bayesian updating. Under the variational formulation of Bayes' rule, streaming RegBayes [167] can naturally be defined as solving:

$$\min_{q(\Theta) \in \mathcal{P}} \text{KL}(q(\Theta)||p_t(\Theta)) + c \cdot \Omega(q(\Theta); B_{t+1}), \quad (27)$$

whose streaming update rule can be derived via convex analysis under a quite general setting.

The sequential updating procedure is perfectly suitable for online learning with data streams, where a revisit to each data point is not allowed. However,

one challenge remains on evaluating the posteriors. If the prior is conjugate to the likelihood model (e.g., a linear Gaussian state-space model) or the state space is discrete (e.g., hidden Markov models [152], [165]), then the sequential updating rule can be done analytically, for example, Kalman filters [97]. In contrast, many complex Bayesian models (e.g., the models involving non-Gaussianity, non-linearity and high-dimensionality) do not have closed-form expression of the posteriors. Therefore, it is computationally intractable to do the sequential update.

4.2.1 Streaming Variational Methods

Various effects have been made to develop streaming variational Bayesian (SVB) methods [43]. Specifically, let \mathcal{A} be a variational algorithm that calculates the approximate posterior q : $q(\Theta) = \mathcal{A}(p(\Theta); B)$. Then, setting $q_0(\Theta) = p_0(\Theta)$, one way to recursively compute an approximation to the posterior is

$$p(\Theta|B_1, \dots, B_{t+1}) \approx q_{t+1}(\Theta) = \mathcal{A}(q_t(\Theta), B_{t+1}). \quad (28)$$

Under the exponential family assumption of q , the streaming update rule has some analytical form.

The streaming RegBayes [167] provides a Bayesian generalization of online passive-aggressive (PA) learning [51], when the posterior regularization term is defined via the max-margin principle. The resulting online Bayesian passive-aggressive (BayesPA) learning adopts a similar streaming variational update to learn max-margin classifiers (e.g., SVMs) in the presence of latent structures (e.g., latent topic representations). Compared to the ordinary PA, BayesPA is more flexible on modeling complex data. For example, BayesPA can discover latent structures via a hierarchical Bayesian treatment as well as allowing for nonparametric Bayesian inference to resolve the complexity of latent components (e.g., using a HDP topic model to resolve the unknown number of topics).

4.2.2 Streaming Monte Carlo Methods

Sequential Monte Carlo (SMC) methods [14], [114], [18] provide simulation-based methods to approximate the posteriors for online Bayesian inference. SMC methods rely on resampling and propagating samples over time with a large number of particles. A standard SMC method would require the full data to be stored for expensive particle rejuvenation to protect particles against degeneracy, leading to an increased storage and processing bottleneck as more data are accrued. For simple conjugate models, such as linear Gaussian state-space models, efficient updating equations can be derived using methods like Kalman filters. For a broader class of models, assumed density filtering (ADF) [107], [142] was developed to extend the computational tractability. Basically, ADF approximates the posterior distribution with a simple conjugate family, leading to approximate

online posterior tracking. Recent improvements on SMC methods include the conditional density filtering (C-DF) method [82], which extends Gibbs sampling to streaming data. C-DF sequentially draws samples from an approximate posterior distribution conditioned on surrogate conditional sufficient statistics, which are approximations to the conditional sufficient statistics using sequential samples or point estimates for parameters along with the data. C-DF requires only data at the current time and produces a provably good approximation to the target posterior.

4.3 Distributed Algorithms

Recent progress has been made on both distributed variational and distributed Monte Carlo methods.

4.3.1 Distributed Variational Methods

If the variational distribution is in some parametric family (e.g., the exponential family), the variational problem can be solved with generic optimization methods. Therefore, the broad literature on distributed optimization [39] provides rich tools for distributed variational inference. However, the disadvantage of a generic solver is that it may fail to explore the structure of Bayesian inference.

First, many Bayesian models have a nature hierarchy, which encodes rich conditional independence structures that can be explored for efficient algorithms, e.g., the distributed variational algorithm for LDA [202]. Second, the inference procedure with Bayes' rule is intrinsically parallelizable. Suppose the data \mathcal{D} is split into non-overlapping batches (often called shards), B_1, \dots, B_M . Then, the Bayes posterior $p(\Theta|\mathcal{D}) = \frac{p_0(\Theta) \prod_{i=1}^M p(B_i|\Theta)}{p(\mathcal{D})}$ can be expressed as

$$p(\Theta|\mathcal{D}) = \frac{1}{C} \prod_{i=1}^M \frac{p_0(\Theta)^{\frac{1}{M}} p(B_i|\Theta)}{p(B_i)} = \frac{1}{C} \prod_{i=1}^M p(\Theta|B_i), \quad (29)$$

where $C = \frac{p(\mathcal{D})}{\prod_{i=1}^M p(B_i)}$. Now, the question is how to calculate the local posteriors (or subset posteriors) $p(\Theta|B_i)$ as well as the normalization factor. The work [43] explores this idea and presents a distributed variational Bayesian method, which approximates the local posterior with an algorithm \mathcal{A} , that is, $p(\Theta|B_i) \approx \mathcal{A}(p_0(\Theta)^{1/M}, B_i)$. Under the exponential family assumption of the prior and the approximate local posteriors, the global posterior can be (approximately) calculated via density product. However, the parametric assumptions may not be reasonable, and the mean-field assumptions can get the marginal distributions right but not the joint distribution.

4.3.2 Distributed Monte Carlo Methods

For MC methods, if independent samples can be directly drawn from the posterior or some proposals (e.g., by using importance sampling), it will be

straightforward to parallelize, e.g., by running multiple independent samplers on separate machines and then simply aggregating the samples [196]. We consider the more challenging cases, where directly sampling from the posterior is intractable and MCMC methods are among the natural choices. There are two groups of methods. One is to run multiple MCMC chains in parallel, and the other is to parallelize a single MCMC chain. The "multiple-chain" parallelism is relatively straightforward if each single chain can be efficiently carried out and an appropriate combination strategy is adopted [68], [196]. However, in Big data applications a single Markov chain itself is often prohibitively slow to converge, due to the massive data sizes or extremely high-dimensional sample spaces. Below, we focus on the methods that parallelize a single Markov chain, for which previous work can be grouped in three major categories.

Blocking: Methods in the first category let each computing unit (e.g., a CPU processor or a GPU core) to perform a part of the computation at each iteration. For example, they independently evaluate the likelihood for each shard across multiple units and combine the local likelihoods with the prior on a master unit to get estimates of the global posterior [172]. Another example is that each computing unit is responsible for updating a part of the state space [192]. These methods involve extensive communications and being problem specific. In these approaches several computing units collaborate to obtain a draw from the posterior.

In order to effectively split the likelihood evaluation or the state space update over multiple computing units, it is important to explore the conditional independence (CI) structure of the model. Many hierarchical Bayesian models naturally have the CI structure (e.g., topic models), while some other models need some transformation to introduce CI structures that are appropriate for parallelization [195].

Divide-and-Conquer: Methods in the second category avoid extensive communication among machines by running independent MCMC chains for each mini-batch and obtaining samples drawn from local posteriors via a single communication. One key step is on combining the samples from local posteriors, which has attracted a lot of attention recently. For example, the Consensus Monte Carlo [164] directly combines local samples by a weighted average, which is valid under an implicit Gaussian assumption while lacking of guarantees for non-Gaussian cases; [137] approximates each local posterior with either an explicit Gaussian or a Gaussian-kernel KDE so that combination follows an explicit density product; [185] builds upon the KDE idea one step further by representing the discrete KDE as a continuous Weierstrass transform; and [127] proposes to calculate the geometric median of local posteriors (or M-posterior), which is provably robust to the presence




Fig. 5. The possible outcomes in two iterations of a Metropolis-Hastings sampler.

of outliers. The M-posterior is approximately solved by the Weiszfeld’s algorithm [26] by embedding the local posteriors in a reproducing kernel Hilbert space.

The potential drawback of these embarrassingly parallel MCMC sampling is that if the local posteriors differ significantly, perhaps due to noise or non-random partitioning of the dataset across nodes, the final combination stage can result in inaccurate global posterior. The recent work [198] presents a context aware distributed Bayesian posterior sampling method to improve inference quality. By allowing nodes to effectively and efficiently share information with each other, each node will eventually draw samples from a more accurate approximate full posterior, and therefore no longer needs any combination.

Prefetching: The idea of prefetching is to make use of parallel processing to calculate multiple likelihoods ahead of time, and only use the ones which are needed. Consider a generic random-walk metropolis-Hastings algorithm at time t . The subsequent steps can be represented by a binary tree, where at each iteration a single new proposal is drawn from a proposal distribution and stochastically accepted or rejected. So, at time $t+n$ the chain has 2^n possible future states, as illustrated in Fig. 5. The vanilla version of prefetching speculatively evaluates all paths in this binary tree [42]. Since only one path of these will be taken, with M cores, this approach achieves a speedup of $\log_2 M$ with respect to single core execution, ignoring communication overheads. More efficient prefetching approaches have been proposed in [16] and [171] by better guessing the probabilities of exploration of both the acceptance and the rejection branches at each node. The recent work [20] presents a delayed acceptance strategy for MH testing, which can be used to improve the efficiency of prefetching.

As a special type of MCMC, Gibbs sampling methods naturally follow a blocking scheme by iterating over some partition of the variables. The early Synchronous Gibbs sampler [69] is highly parallel by sampling all variables simultaneously on separate processors. However, the extreme parallelism comes at a cost, e.g., the sampler may not converge to the correct stationary distribution in some cases [77]. The work [77] develops various variable partitioning strategies to achieve fast parallelization, while maintaining the convergence to the target posterior, and the work [93] analyzes the convergence and correctness of

means of the Synchronous Gibbs sampler (a.k.a, the Hogwild parallel Gibbs sampler) for sampling from Gaussian distributions. Many other parallel Gibbs sampling algorithms have been developed for specific models. For example, various distributed Gibbs samplers [138], [168], [8], [115] have been developed for the vanilla LDA, [48] develops a distributed Gibbs sampler via data augmentation to learn large-scale topic graphs with a logistic-normal topic model, and parallel algorithms for DP mixtures have been developed by introducing auxiliary variables for additional CI structures [195], while with the potential risk of causing extremely imbalanced partitions [66].

Note that the stochastic methods and distributed computing are not exclusive. Combining both often leads to more efficient solutions. For example, for optimization methods, parallel SGD methods have been extensively studied [212], [140]. In particular, [140] presents a parallel SGD algorithm without locks, called Hogwild!, where multiple processors are allowed equal access to the shared memory and are able to update individual components of memory at will. Such a scheme is particularly suitable for sparse learning problems. For Bayesian methods, the distributed stochastic gradient Langevin dynamics (D-SGLD) method has been developed in [10].

5 TOOLS, SOFTWARE AND SYSTEMS

Though stochastic algorithms are easy to implement, distributed methods often need a careful design of the system architectures and programming libraries. For system architectures, we may have a shared memory computer with many cores, a cluster with many machines interconnected by network (either commodity or high-speed), or accelerating hardware like graphics processing units (GPUs), field-programmable gate arrays (FPGAs), and application specific integrated circuits (ASICs). In this section, we review the distributed programming libraries suitable for various system architectures and summarize the existing tools for Bayesian inference.

5.1 System Primitives

Every architecture has its low-level libraries, in which the architectures (e.g., threads, machines, or GPU cores) are visible to the programmer.

Shared Memory Computer: A shared memory computer passes data from one CPU core to another by simply storing it into the main memory. Therefore, the communication latency is low. It is also easy to program and acquire. Meanwhile it is prevalent—it is the basic component of large distributed clusters and host of GPUs or other accelerating hardware. Due to these reasons, writing a multi-thread program is usually the first step towards large-scale learning. However, its drawbacks include limited memory/IO

capacity and bandwidth, and restricted scalability, which can be addressed by distributed clusters.

Programmers work with threads in a shared memory setting. A threading library supports: 1) spawning a thread and wait it to complete; 2) synchronization: method to prevent conflict access of resources, such as locks; 3) atomic: operations, such as increment that can be executed in parallel safely. Besides threads and locks, there are alternative programming frameworks. For example, Scala uses *actor*, which responds to a message that it receives; Go uses *channel*, which is a multi-provider, multi-consumer queue. There are also libraries automating specific parallel pattern, e.g., OpenMP [6] supports parallel patterns like parallel for or reduction, and synchronization patterns like barrier; TBB [4] has pipeline and lightweight green threads. Choosing right programming models sometimes can simplify the implementation.

Accelerating Hardware: GPUs are self-contained parallel computational devices that can be housed in conventional desktop or laptop computers. A single GPU can provide floating operations per second (FLOPS) performance as good as a small cluster. Yet compared to conventional multi-core processors, GPUs are cheap, easily accessible, easy to maintain, easy to code, and dedicated local devices with low power consumption. GPUs follow a single instruction multiple data (SIMD) pattern, i.e., a single program will be executed on all cores given different data. This pattern is suitable for many ML applications. However, GPUs may be limited due to: 1) small memory capacity; 2) restricted SIMD programming model; and 3) high CPU-GPU or GPU-GPU communication latency.

CUDA is a framework for Nvidia GPUs. In CUDA, user writes *kernels*, which is just a function. Then the *one* kernel, will be run *all* cores of a user defined set of cores. Users can also specify the location (CPU/GPU global memory or GPU blockwise memory) of data and move them around. OpenCL is another framework, which supports Nvidia, AMD, and Intel GPUs. There are also higher level libraries such as OpenACC, which is the GPU counterpart of OpenMP.

Many Bayesian inference methods have been accelerated with GPUs. For example, [172] adopts GPUs to parallelize the likelihood evaluation in MCMC; [109] provides GPU parallelization for population-based MCMC methods [89] as well as SMC samplers [131]; and [25] uses GPU computing to develop fast Hamiltonian Monte Carlo methods. For variational Bayesian methods, [199] demonstrates an example of using GPUs to accelerate the collapsed variational Bayesian algorithm for LDA. BIDMach [46] is a distributed GPU framework for machine learning. In particular, BIDMach LDA with a single GPU is able to learn faster than the state-of-the-art CPU based LDA implementation [8], which use 100 CPUs.

Finally, acceleration with other hardware (e.g., FP-




Fig. 6. Various architectures: (a) MapReduce/Spark; (b) Pregel/GraphLab; (c) Parameter servers.

GAs) has also been investigated [47].

Distributed Cluster: For distributed clusters, a low-level framework should allow users to do: 1) *Communication*: sending and receiving data from/to another machine or a group of machines; 2) *Synchronization*: synchronize the processes; 3) *Fault handling*: decide what to do if a process/machine breaks down. For example, MPI provides a set of primitives including *send*, *receive*, *broadcast* and *reduce* for communication. MPI also provides synchronization operations, such as *barrier*. MPI handles fault by simply terminating all processes. MPI works on various network infrastructures, such as ethernet or Infiniband. Besides MPI, there are other frameworks that support communication, synchronization and fault handling, such as 1) message queues, where processes can put and get messages from globally shared message queues; 2) remote procedural calls (RPCs), where a process can invoke a procedure in another process, passing its own data to that remote procedure, and finally get execution results. MrBayes [159], [12] provides a MPI-based parallel algorithm for Metropolis-coupled MCMC for Bayesian phylogenetic inference.

Programming with system primitive libraries are most flexible and lightweight. However for sophisticated applications, which may require asynchronous execution, need to modify the global parameters while running, or need many parallel execution blocks, it would be painful and error prone to write the parallel code using the low-level system primitives. Below, we review some high-level distributed computing frameworks, which automatically execute the user declared tasks on desired architectures. We refer the readers to the book [27] for more details on GPUs, MapReduce, and some other large-scale machine learning examples (e.g., using parallel online learning).

5.2 MapReduce and Spark

MapReduce [55] is a distributed computing framework for key-value stores. It reads key-value stores from disk, performs some transformations to these key-value stores in parallel, and writes the final results to disk. A typical MapReduce cycle is as follows:

- 1) Spawn some workers on all machines;
- 2) Workers read input key-value pairs in parallel from a distributed file system;
- 3) *Map*: Pass each key-value pair to a user defined function, which will generate some intermediate key-value pairs;

- 4) According to the key, hash the intermediate key-value pairs to all machines, then merge key-value pairs that have the same key, result with (key, list of values) pairs;
- 5) *Reduce*: In parallel, pass each (key, list of values) pairs to a user defined function, which will generate some output key-value pairs;
- 6) Write output key-value pairs to the file system.

There are two user defined functions, *mapper* and *reducer*. For ML, a key-value store is often data samples, mapper is often used for computing latent variables, likelihoods or gradients for each data sample, and reducer is often used to aggregate the information from each data sample, where the information can be used for estimating parameters or checking convergence. [50] discusses a number of ML algorithms on MapReduce, including linear regression, naive Bayes, neural networks, PCA, SVM, etc. Mahout [1] is a ML package built upon Hadoop, an open source implementation of MapReduce. Mahout provides collaborative filtering, classification, clustering, dimensionality reduction and topic modeling algorithms. [202] is a MapReduce based LDA. However, a major drawback of MapReduce is that it needs to read the data from disk at *every iteration*. The overhead of reading data becomes dominant for many iterative ML algorithms as well as interactive data analysis tools [201].

Spark [201] is another framework for distributed ML methods that involve iterative jobs. The core of Spark is resilient data sets (RDDs), which is essentially a dataset distributed across machines. RDD can be stored either in memory or disk: Spark decides it automatically, and users can provide hints to Spark which to store in memory. This avoids reading the dataset at every iteration. Users can perform *parallel operations* to RDDs, which will transform a RDD to another. Available parallel operations are like *foreach* and *reduce*. We can use *foreach* to do the computation for each data, and use *reduce* to aggregate information from data. Because parallel operations are just a parallel version of the corresponding serial operations, a Spark program looks almost identical to its serial counterpart. Spark can outperform Hadoop for iterative ML jobs by 10x, and is able to interactively query a 39GB dataset in 1 second [201].

5.3 Iterative Graph Computing

Both MapReduce and Spark have a star architecture as in Fig. 6 (a), where only master-slave communication is permitted; they do not allow one key-value pair to interact with another, e.g., reading or modifying the value of another key-value pair. The interaction is necessary for applications like PageRank, Gibbs sampling, and variational Bayes optimized by coordinate descent, all of which require variables to get their own values based on other related variables. Hence there comes graph computing, where the computational

task is defined by a sparse graph that specifies the data dependency, as shown in Fig. 6 (b).

Pregel [120] is a bulk synchronous parallel (BSP) graph computing engine. The computation model is a sparse graph with data on vertices and edges, where each vertex receives all messages sent to it in the last iteration; updates data on the vertex based on the messages; and sends out messages along adjacent edges. For example, Gibbs sampling can be done easily by sending the vertex statistics to adjacent vertices and then the conditional probability can be computed. GPS [162] is an open source implementation of Pregel with new features (e.g., dynamic graph repartition).

GraphLab [117] is a more sophisticated graph computing engine that allows asynchronous execution and flexible scheduling. A GraphLab iteration picks up a vertex v in the task queue; and passes the vertex to a user defined function, which may modify the data on the vertex, its adjacent edges and vertices, and finally may add its adjacent vertices to the task queue. Note that several nodes can be evaluated in parallel as long as they do not violate the consistency guarantee which ensures that GraphLab is equivalent with some serial algorithm. It has been used to parallelize a number of ML tasks, including matrix factorization, Gibbs sampling and Lasso [117]. [210] presents a distributed Gibbs sampler on GraphLab for an improved sLDA model using RegBayes. Several other graph computing engines have been developed. For example, GraphX [197] is an extension of Spark for graph computing; and GraphChi [106] is a disk based version of GraphLab.

5.4 Parameter Servers

All the above frameworks restrict the communication between workers. For example, MapReduce and Spark don't allow communication between workers, while Pregel and GraphLab only allow vertices to communicate with adjacent nodes. On the other side, many ML methods follow a pattern that

- Data are partitioned on many workers;
- There are some shared global parameters (e.g., the model weights in a gradient descent method or the topic-word count matrix in the collapsed Gibbs sampler for LDA [81]);
- Workers fetch data and update (parts of) global parameters based on their local data (e.g., using the local gradients or local sufficient statistics).

Though it is straightforward to implement on shared memory computers, it is rather difficult in a distributed setting. The goal of parameter servers is to provide a distributed data structure for parameters.

A parameter server is a key-value store (like a hash map), accessible for all workers. It supports for *get* and *set* (or *update*) for each entry. In a distributed setting, both server and client consist of many nodes (see Fig. 6 (c)). Memcached [3] is an

in memory key-value store that provides `get` and `set` for arbitrary data. However it doesn't have a mechanism to resolve conflicts raised by concurrent access, e.g. concurrent writes for a single entry. Applications like [168] require to lock the global entry while updating, which leads to suboptimal performance. Piccolo [151] addresses this by introducing user-defined accumulations, which correctly address concurrent updates to the same key. Piccolo has a set of built-in user defined accumulations such as summation, multiplication, and min/max.

One important tradeoff made by parameter servers is that they sacrifice consistency for less latency—`get` may not return the most recent value, so that it can return immediately without waiting for most recent updates to reach the server. While this improves the performance significantly, it can potentially slow down convergence due to outdated parameters. [86] proposed Stale Synchronous Parallel (SSP), where the staleness of parameters is *bounded* and the fastest worker can be ahead of the slowest one by no more than τ iterations, where τ can be tuned to get a fast convergence as well as low waiting time. Petuum [52] is a SSP based parameter server. [113] proposed communication-reducing improvements, including key caching, message compression and message filtering, and it also supports elastically adding and removing both server and worker nodes.

Parameter servers have been deployed in learning very large-scale logistic regression [113], deep networks [54], LDA [8], [111] and Lasso [52]. [113] learns a 2000-topic LDA with 5 billion documents and 5 million unique tokens on 6000 machines in 20 hours. Yahoo! LDA [8] has a parameter server designed specifically for Bayesian latent variable models and it is the fastest available LDA software. There are a bunch of distributed topic modeling softwares based on Yahoo! LDA, including [211] for MedLDA and [48] for correlated topic models.

5.5 Model Parallel Inference

MapReduce, Spark and Parameter servers take the *data-parallelism* approach, where *data* are partitioned across machines and computations are performed on each node given a copy of the globally shared *model*. However, as the model size rapidly grows (i.e., the large M challenge), the models cannot fit in a single computer's memory. Model-parallelism addresses this challenge by partitioning the *model* and storing a part of the model on each node. Then, partial updates (i.e., the updates of model parts) are carried out on each node. Benefits of model-parallelism include large model sizes, flexibility to focus workers on fastest-converging parameters, and more accurate convergence because no delayed update is involved.

STRADS [110] provides primitives for model-parallelism and it handles the distributed storage of model and data automatically. STRADS requires that

a partial update could be computed using just the model part together with data. Users writes *schedule* that assigns model sets to workers, *push* that computes the partial updates for model, *pop* that applies updates to model. An automatic *sync* primitive will ensure that users always get the latest model. As a concrete example, [204] demonstrates a model parallel LDA, in which both data and model are partitioned by vocabulary. In each iteration, a worker only samples latent variables and updates the model related to the vocabulary part assigned to it. The model then rotates between workers, until a full cycle is completed. Unlike data parallel LDA [168], [8], [138], the sampler always uses the latest models and no read-write lock is needed on models, thereby leading to faster convergence than data-parallel LDAs.

Note that model-parallelism is not a replacement but a complement of data-parallelism. For example, [186] showed a two layer LDA system, where layer 1 is model-parallelism and layer 2 consists of several local model-parallelism clusters performing asynchronous updates on an globally distributed model.

6 RELATED WORK

We briefly review some closely related topics on Bayesian methods.

6.1 Sparse Bayesian Methods

Sparsity regularization has been very effective in controlling model complexity as well as identifying important factors with parsimonious estimates [64], [19] when learning in high-dimensional spaces. Though such regularized estimates can be viewed as finding the maximum a posteriori (MAP) estimates of a Bayesian model, they are not truly Bayesian. A Bayes estimator optimizes the Bayes risk, which is the expectation of a loss averaged over the posterior distribution. Furthermore, a Bayesian approach takes the uncertainty into consideration by inferring the entire posterior distribution, rather than a single point. Popular sparse Bayesian methods include those using spike-and-slab priors [129], [88], those using adaptive shrinkage with heavily-tailed priors (e.g., a Laplace prior) [146], and the methods using model space search. We refer the readers to [141] for a nice review. The recent work [133] presents a Bayesian variable selection method with strong selection consistency.

6.2 Bayesian Optimization

Bayesian optimization (BO) [41] aims to optimize some objective, which may be expensive to evaluate or do not have easily available derivatives, with successful applications in robotics, planning, recommendation, advertising, and automatic algorithm configuration. In Big data era, the learning models are becoming incredibly huge [108] and the learning algorithms often have tuning parameters (e.g., the

learning rates of SGD algorithms). Manually tuning such hyper-parameters is often prohibitive. Recent progress has been made on practical BO methods to automatically select good parameters based on Gaussian process using a multi-core parallel Monte Carlo algorithm [169] or a stochastic variational inference method [139]. To deal with the challenge of learning in high dimensional spaces, the work [188] presents a random embedding BO algorithm.

6.3 Small Variance Asymptotics

Small variance asymptotics (SVA) conjoins Bayesian nonparametrics with optimization. It setups conceptual links between probabilistic and non-probabilistic models and derives new algorithms that can be simple and scalable. For example, a Gaussian mixture model (GMMs) reduces to K-means when the likelihood variance goes to zero; and similarly the probabilistic PCA reduces to a standard PCA by letting the covariance of the likelihood in pPCA approach zero [179]. Recent progress has been made on deriving new computational methods by performing SVA analysis to Bayesian nonparametric models. For example, DP-means is an extension of K-means for nonparametric inference by doing SVA analysis to DP mixtures [105], [92]. Similar analysis has been done for hidden Markov models [160], latent feature models [44], and DP mixtures of SVMs [187], [208] which perform clustering and classification in a joint framework. Note that the progress on small-variance techniques is orthogonal to the advances in big learning. For example, DP-Means has been scaled to deal with massive data using distributed computing [144].

7 CONCLUSIONS AND DISCUSSIONS

We present a survey of recent advances on big learning with Bayesian methods, including Bayesian nonparametrics, regularized Bayesian inference, and scale inference algorithms and systems based stochastic subsampling or distributed computing. It is helpful to note that our review is not exhaustive. Big learning has attracted intense interest with active research spanning diverse fields, including machine learning, databases, parallel and distributed systems, and programming languages.

Current ML methods still require considerable human expertise in devising appropriate features, priors, and models. To make ML more widely used and eventually become a common part of our day to day tools in data sciences, several promising projects have been started. Google prediction API is one of the earliest efforts that try to make ML accessible for beginners by providing easy-to-use cloud service. Microsoft AzureML takes a similar approach by providing a visual interface to help design experiments. SystemML [76] provides an R-like declarative language to specify ML tasks based on MapReduce,

and MLBase [104] further improves it by providing learning-specific optimizer that transforms a declarative task into a sophisticated learning plan. Finally, Automated Statistician (AutoStat) [116] aims to automate the process of statistical modeling, by using Bayesian model selection strategies to automatically choose good models/features and to interpret the results in easy-to-understand ways, in terms of automatically generated reports. Such efforts would have a tremendous impact on fields that currently rely on expert statisticians, ML researchers, and data scientists.

ACKNOWLEDGEMENTS

The work is supported by National 973 Projects (2013CB329403, 2012CB316301), NSF of China Projects (61322308, 61332007), and Tsinghua Initiative Scientific Research Program (20121088071).

REFERENCES

- [1] Apache mahout: <https://mahout.apache.org/>.
- [2] <http://lshtc.iit.demokritos.gr/>.
- [3] <http://memcached.org>.
- [4] <https://www.threadingbuildingblocks.org/>.
- [5] <http://www.image-net.org/about-overview>.
- [6] <http://www.openmp.org>.
- [7] R. Adams, H. Wallach, and Z. Ghahramani. Learning the structure of deep sparse graphical models. In *AISTATS*, 2010.
- [8] A. Ahmed, M. Aly, J. Gonzalez, S. Narayanamurthy, and A. Smola. Scalable inference in latent variable models. In *WSDM*, 2012.
- [9] S. Ahn, A. Korattikara, and M. Welling. Bayesian posterior sampling via stochastic gradient fisher scoring. In *ICML*, 2012.
- [10] S. Ahn, B. Shahbaba, and M. Welling. Distributed stochastic gradient MCMC. In *ICML*, 2014.
- [11] J. Aitchison and S. M. Shen. Logistic-normal distributions: Some properties and uses. *Biometrika*, 67(2):261–272, 1980.
- [12] G. Altekars, S. Dwarkadas, J. Huelsenbeck, and F. Ronquist. Parallel Metropolis coupled Markov chain Monte Carlo for Bayesian phylogenetic inference. *Bioinformatics*, 20(3):407–415, 2004.
- [13] S. Amari. Natural gradient works efficiently in learning. *Neural Comput.*, 10:251–276, 1998.
- [14] C. Andrieu, A. Doucet, and R. Holenstein. Particle Markov chain Monte Carlo methods. *J. R. Stat. Soc., Ser. B*, 72(3):269–342, 2010.
- [15] C. Andrieu, N. D. Freitas, A. Doucet, and M. I. Jordan. An introduction to MCMC for machine learning. *Machine Learning*, 50:5–43, 2003.
- [16] E. Angelino, E. Kohler, A. Waterland, M. Seltzer, and R. P. Adams. Accelerating MCMC via parallel predictive prefetching. *arXiv:1403.7265*, 2014.
- [17] C. Antoniak. Mixture of Dirichlet process with applications to Bayesian nonparametric problems. *Ann. Stats.*, (273):1152–1174, 1974.
- [18] M. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for online nonlinear/non-gaussian Bayesian tracking. *IEEE Trans. Signal Process.*, 50(2):174–188, 2002.
- [19] F. Bach, R. Jenatton, J. Mairal, and G. Obozinski. Optimization with sparsity-inducing penalties. *Foundations and Trends in Machine Learning*, 4(1):1–106, 2012.
- [20] M. Banterle, C. Grazian, and C. P. Robert. Accelerating Metropolis-Hastings algorithms: Delayed acceptance with prefetching. *arXiv:1406.2660*, 2014.
- [21] R. Bardenet, A. Doucet, and C. Holmes. Towards scaling up Markov chain Monte Carlo: an adaptive subsampling approach. In *ICML*, 2014.

- [22] R. Bardenet and O.-A. Maillard. Concentration inequalities for sampling without replacement. *arXiv:1309.4029*, 2013.
- [23] M. Beal, Z. Ghahramani, and C. Rasmussen. The infinite hidden markov model. In *NIPS*, 2002.
- [24] M. J. Beal. Variational algorithms for approximate Bayesian inference. *PhD Thesis, University of Cambridge*, 2003.
- [25] A. L. Beam, S. K. Ghosh, and J. Doyle. Fast Hamiltonian Monte Carlo using GPU computing. *arXiv:1402.4089*, 2014.
- [26] A. Beck and S. Sabach. Weiszfelds method: Old and new results. *J. of Opt. Theory and Applications*, 2014.
- [27] R. Bekkerman, M. Bilenko, and J. Langford. *Scaling up machine learning: Parallel and distributed approaches*. Cambridge University Press, 2011.
- [28] S. Bengio, J. Weston, and D. Grangier. Label embedding trees for large multi-class tasks. In *NIPS*, 2010.
- [29] Y. Bengio, A. Courville, and P. Vincent. Representation Learning: A Review and New Perspectives. *IEEE Trans. on PAMI*, 35(8):1798–1828, 2013.
- [30] W. Bialek, I. Nemenman, and N. Tishby. Predictability, complexity and learning. *Neural Comput.*, 13:2409–2463, 2001.
- [31] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- [32] D. Blei and P. Frazier. Distance dependent Chinese restaurant processes. In *ICML*, 2010.
- [33] D. Blei and M. Jordan. Variational inference for dirichlet process mixtures. *Bayesian Analysis*, 1:121–144, 2006.
- [34] D. Blei and J. Lafferty. Correlated topic models. In *NIPS*, 2006.
- [35] D. Blei and J. McAuliffe. Supervised topic models. In *NIPS*, 2007.
- [36] D. Blei, A. Ng, and M. I. Jordan. Latent Dirichlet Allocation. *JMLR*, (3):993–1022, 2003.
- [37] L. Bottou. *Online Algorithms and Stochastic Approximations*. Online Learning and Neural Networks, Edited by David Saad, Cambridge University Press, Cambridge, UK, 1998.
- [38] L. Bottou and O. Bousquet. The tradeoffs of large scale learning. In *NIPS*, 2008.
- [39] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. *Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers*, volume 3. Foundations and Trends in Machine Learning, 2011.
- [40] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [41] E. Brochu, V. M. Cora, and N. de Freitas. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv:1012.2599*, 2010.
- [42] A. Brockwell. Parallel Markov chain Monte Carlo simulation by pre-fetching. *JCGS*, 15(1):246–261, 2006.
- [43] T. Broderick, N. Boyd, A. Wibisono, A. C. Wilson, and M. I. Jordan. Streaming variational Bayes. In *NIPS*, 2013.
- [44] T. Broderick, B. Kulis, and M. I. Jordan. MAD-Bayes: MAP-based asymptotic derivations from Bayes. In *ICML*, 2013.
- [45] G. Brumfiel. High-energy physics: Down the petabyte highway. *Nature*, 469:282–283, 2011.
- [46] J. Canny and H. Zhao. Bidmach: Large-scale learning with zero memory allocation. In *NIPS workshop on "Big Learning"*, 2013.
- [47] T. Chau, J. Targett, M. Wijeyasinghe, W. Luk, P. Cheung, B. Cope, A. Eele, and J. Maciejowski. Accelerating sequential Monte Carlo method for real-time air traffic management. *SIGARCH Computer Architecture News*, 41(5):35–40, 2013.
- [48] J. Chen, J. Zhu, Z. Wang, X. Zheng, and B. Zhang. Scalable inference for logistic-normal topic models. In *NIPS*, 2013.
- [49] T. Chen, E. B. Fox, and C. Guestrin. Stochastic gradient Hamiltonian Monte Carlo. In *ICML*, 2014.
- [50] C. Chu, S. K. Kim, Y.-A. Lin, Y. Yu, G. Bradski, A. Y. Ng, and K. Olukotun. Map-reduce for machine learning on multicore. In *NIPS*, 2007.
- [51] K. Crammer, O. Dekel, J. Keshet, S. Shalev-Shwartz, and Y. Singer. Online passive-aggressive algorithms. *JMLR*, (7):551–585, 2006.
- [52] W. Dai, J. Wei, X. Zheng, J. K. Kim, S. Lee, J. Yin, Q. Ho, and E. Xing. Petuum: A framework for iterative-convergent distributed ML. In *arXiv:1312.7651*, 2013.
- [53] P. Dallaire, P. Giguere, and B. Chaib-draa. Learning the structure of probabilistic graphical models with an extended cascading Indian buffet process. In *AAAI*, 2014.
- [54] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, M. Mao, A. Senior, P. Tucker, K. Yang, Q. V. Le, et al. Large scale distributed deep networks. In *NIPS*, 2012.
- [55] J. Dean and S. Ghemawat. MapReduce: Simplified data processing on large clusters. In *OSDI*, 2004.
- [56] J. Deng, S. Satheesh, A. Berg, and L. Fei-Fei. Fast and balanced: Efficient label tree learning for large scale object recognition. In *NIPS*, 2011.
- [57] C. Doctorow. Big data: Welcome to the petacentre. *Nature*, 455:16–21, 2008.
- [58] S. Donnet, V. Rivoirard, J. Rousseau, and C. Scricciolo. On convergence rates of empirical Bayes procedures. In *47th Scientific Meeting of the Italian Statistical Society*, 2014.
- [59] F. Doshi-Velez, K. Miller, J. V. Gael, and Y. W. Teh. Variational inference for the Indian buffet process. In *AISTATS*, 2009.
- [60] J. Duan, M. Guindani, and A. Gelfand. Generalized spatial Dirichlet process models. *Biometrika*, 94(4), 2007.
- [61] J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *JMLR*, 12:2121–2159, 2011.
- [62] D. V. Dyk and X. Meng. The art of data augmentation. *JCGS*, 10(1):1–50, 2001.
- [63] B. Efron. Bayes’ theorem in the 21st century. *Science*, 340(6137):1177–1178, 2013.
- [64] J. Fan and J. Lv. A selective overview of variable selection in high dimensional feature space. *Statist. Sinica*, 20:101–148, 2010.
- [65] T. Ferguson. A Bayesian analysis of some nonparametric problems. *Ann. Stats.*, (1):209–230, 1973.
- [66] Y. Gal and Z. Ghahramani. Pitfalls in the use of parallel inference for the Dirichlet process. In *ICML*, 2014.
- [67] A. Gelman, J. Carlin, H. Stern, D. Dunson, A. Vehtari, and D. Rubin. *Bayesian Data Analysis*. Third Edition, Chapman & Hall/CRC Texts in Statistical Science, 2013.
- [68] A. Gelman and D. Rubin. Inference from iterative simulation using multiple simulations. *Statist. Sci.*, 7(4):457–511, 1992.
- [69] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. on PAMI*, 6(1):721–741, 1984.
- [70] E. I. George and D. P. Foster. Calibration and empirical Bayes variable selection. *Biometrika*, 87(4):731–747, 2000.
- [71] S. Gershman, P. Frazier, and D. Blei. Distance dependent infinite latent feature models. *arXiv:1110.5454*, 2011.
- [72] S. Gershman and D. Blei. A tutorial on Bayesian nonparametric models. *J. Math. Psychol.*, (56):1–12, 2012.
- [73] C. J. Geyer and E. A. Thompson. Annealing Markov chain Monte Carlo with applications to ancestral inference. *JASA*, 90(431):909–920, 1995.
- [74] Z. Ghahramani. Bayesian nonparametrics and the probabilistic approach to modelling. *Phil. Trans. of the Royal Society*, 2013.
- [75] J. K. Ghosh and R. Ramamoorthi. *Bayesian Nonparametrics*. Springer, New York, NY, 2003.
- [76] A. Ghoting, R. Krishnamurthy, E. Pednault, B. Reinwald, V. Sindhwani, and et al. SystemML: Declarative machine learning on MapReduce. In *ICDE*, 2011.
- [77] J. E. Gonzalez, Y. Low, A. Gretton, and C. Guestrin. Parallel Gibbs sampling: From colored fields to thin junction trees. In *AISTATS*, 2011.
- [78] P. Gopalan and D. Blei. Efficient discovery of overlapping communities in massive networks. *PNAS*, 110(36):14534–14539, 2013.
- [79] A. Grelaud, C. P. Robert, J.-M. Marin, F. Rodolphe, and J.-F. Taly. Likelihood-free methods for model choice in Gibbs random fields. *Bayesian Analysis*, 4(2):317–336, 2009.
- [80] T. Griffiths and Z. Ghahramani. Infinite latent feature models and the Indian buffet process. In *NIPS*, 2006.
- [81] T. Griffiths and M. Steyvers. Finding scientific topics. *PNAS*, 2004.
- [82] R. Guhaniyogi, S. Qamar, and D. Dunson. Bayesian conditional density filtering for big data. *arXiv:1401.3632*, 2014.
- [83] W. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.

- [84] G. Hinton, L. Deng, D. Yu, A. Mohamed, N. Jaitly, and etc. Deep neural networks for acoustic modeling in speech recognition. *IEEE Signal Process. Mag.*, 29(6):82–97, 2012.
- [85] N. Hjort, C. Holmes, P. Muller, and S. Walker. *Bayesian Nonparametrics: Principles and Practice*. Cambridge University Press, 2010.
- [86] Q. Ho, J. Cipar, H. Cui, S. Lee, J. Kim, P. Gibbons, G. Gibson, G. Ganger, and E. Xing. More effective distributed ML via a stale synchronous parallel parameter server. In *NIPS*, 2013.
- [87] M. D. Hoffman, D. Blei, C. Wang, and J. Paisley. Stochastic variational inference. *JMLR*, 14:1303–1347, 2013.
- [88] H. Ishwaran and J. S. Rao. Spike and slab variable selection: Frequentist and Bayesian strategies. *Ann. Statist.*, 33:730–773, 2005.
- [89] A. Jasra, D. A. Stephens, and C. C. Holmes. On population-based simulation for static inference. *Statistics and Computing*, 17(3):263–279, 2007.
- [90] E. T. Jaynes. Prior probabilities. *IEEE Trans. on Sys. Sci. and Cybernetics*, 4:227–241, 1968.
- [91] H. Jeffreys. An invariant form for the prior probability in estimation problems. *Proc. of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 186(1007):453–461, 1945.
- [92] K. Jiang, B. Kulis, and M. Jordan. Small-variance asymptotics for exponential family Dirichlet process mixture models. In *NIPS*, 2012.
- [93] M. J. Johnson, J. Saunderson, and A. S. Willsky. Analyzing hogwild parallel gaussian Gibbs sampling. In *NIPS*, 2013.
- [94] M. Jordan. The era of big data. *ISBA Bulletin*, 18(2):1–3, 2011.
- [95] M. Jordan, Z. Ghahramani, T. Jaakkola, and L. Saul. An introduction to variational methods for graphical models. *MLJ*, 37(2):183–233, 1999.
- [96] J. B. Kadane and N. A. Lazar. Methods and criteria for model selection. *JASA*, 99(465):279–290, 2004.
- [97] R. E. Kalman. A new approach to linear filtering and prediction problems. *J. Fluids Eng.*, 82(1):35–45, 1960.
- [98] R. E. Kass and A. E. Raftery. Bayes factors. *JASA*, 90(430):773–795, 1995.
- [99] D. I. Kim, P. Gopalan, D. M. Blei, and E. B. Sudderth. Efficient online inference for Bayesian nonparametric relational models. In *NIPS*, 2013.
- [100] D. Kingma and M. Welling. Auto-encoding variational bayes. In *ICLR*, 2014.
- [101] D. Kingma and M. Welling. Efficient gradient-based inference through transformations between Bayes nets and neural nets. In *ICML*, 2014.
- [102] A. Korattikara, Y. Chen, and M. Welling. Austerity in MCMC land: Cutting the Metropolis-Hastings budget. In *ICML*, 2014.
- [103] O. Koyejo and J. Ghosh. Constrained Bayesian inference for low rank multitask learning. In *UAI*, 2013.
- [104] T. Kraska, A. Talwalkar, J. Duchi, R. Griffith, M. Franklin, and M. Jordan. MLbase: A distributed machine-learning system. In *CIDR*, 2013.
- [105] B. Kulis and M. I. Jordan. Revisiting k-means: New algorithms via Bayesian nonparametrics. In *ICML*, 2012.
- [106] A. Kyrola, G. E. Blelloch, and C. Guestrin. GraphChi: Large-scale graph computation on just a PC. In *OSDI*, 2012.
- [107] S. L. Lauritzen. Propagation of probabilities, means and variances in mixed graphical association models. *JASA*, 87:1098–1108, 1992.
- [108] Q. Le, M. Ranzato, R. Monga, M. Devin, K. Chen, G. Corrado, J. Dean, and A. Ng. Building high-level features using large scale unsupervised learning. In *ICML*, 2012.
- [109] A. Lee, C. Yau, M. B. Giles, A. Doucet, and C. C. Holmes. On the utility of graphics cards to perform massively parallel simulation of advanced Monte Carlo methods. *JCGS*, 19(4):769–789, 2010.
- [110] S. Lee, J. K. Kim, X. Zheng, Q. Ho, G. Gibson, and E. Xing. Primitives for dynamic big model parallelism. *arXiv:1406.4580*, 2014.
- [111] A. Q. Li, A. Ahmed, S. Ravi, and A. J. Smola. Reducing the sampling complexity of topic models. In *SIGKDD*, 2014.
- [112] F.-F. Li and P. Perona. A bayesian hierarchical model for learning natural scene categories. In *CVPR*, 2005.
- [113] M. Li, D. Andersen, J. W. Park, A. Smola, A. Ahmed, V. Josifovski, J. Long, E. Shekita, and B.-Y. Su. Scaling distributed machine learning with the parameter server. In *OSDI*, 2014.
- [114] J. S. Liu and R. Chen. Sequential Monte Carlo methods for dynamic systems. *JASA*, 93(443):1032–1044, 1998.
- [115] Z. Liu, Y. Zhang, E. Y. Chang, and M. Sun. PLDA+: Parallel latent Dirichlet allocation with data placement and pipeline processing. *TIST*, 2(3), 2011.
- [116] J. Lloyd, D. Duvenaud, R. Grosse, J. Tenenbaum, and Z. Ghahramani. Automatic construction and natural language description of nonparametric regression models. In *AAAI*, 2014.
- [117] Y. Low, J. Gonzalez, A. Kyrola, D. Bickson, C. Guestrin, and J. Hellerstein. Graphlab: A new framework for parallel machine learning. In *UAI*, 2013.
- [118] S. MacEachern. Dependent nonparametric processes. In *ASA proceedings of the section on Bayesian statistical science*, 1999.
- [119] D. Maclaurin and R. P. Adams. Firefly Monte Carlo: Exact MCMC with subsets of data. In *UAI*, 2014.
- [120] G. Malewicz, M. H. Austern, A. J. Bik, J. C. Dehnert, I. Horn, N. Leiser, and G. Czajkowski. Pregel: a system for large-scale graph processing. In *SIGMOD*, 2010.
- [121] S. Mandt and D. Blei. Smoothed gradients for stochastic variational inference. *arXiv:1406.3650*, 2014.
- [122] B. Marlin, E. Khan, and K. Murphy. Piecewise bounds for estimating Bernoulli-logistic latent Gaussian models. In *ICML*, 2011.
- [123] J. D. McAuliffe, D. M. Blei, and M. I. Jordan. Nonparametric empirical Bayes for the Dirichlet process mixture model. *Statistics and Computing*, 16(1):5–14, 2006.
- [124] S. Mei, J. Zhu, and X. Zhu. Robust RegBayes: Selectively incorporating first-order logic domain knowledge into Bayesian models. In *ICML*, 2014.
- [125] N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller. Equation of state calculations by fast computing machines. *J. Chem. Phys.*, 21(6):1087, 1953.
- [126] K. Miller, T. Griffiths, and M. Jordan. Nonparametric latent feature models for link prediction. In *NIPS*, 2009.
- [127] S. Minsker, S. Srivastava, L. Lin, and D. B. Dunson. Scalable and robust Bayesian inference via the median posterior. In *ICML*, 2014.
- [128] T. Mitchell. *Machine Learning*. McGraw Hill, 1997.
- [129] T. Mitchell and J. Beauchamp. Bayesian variable selection in linear regression. *JASA*, 83(404):1023–1032, 1988.
- [130] A. Mnih and K. Gregor. Neural variational inference and learning in belief networks. In *ICML*, 2014.
- [131] P. D. Moral, A. Doucet, and A. Jasra. Sequential Monte Carlo samplers. *J. R. Stat. Soc., Ser B*, 68(3):411–436, 2006.
- [132] P. Muller and F. A. Quintana. Nonparametric Bayesian data analysis. *Statistical Science*, 19(1):95–110, 2004.
- [133] N. Narisetty and X. He. Bayesian variable selection with shrinking and diffusing priors. *Ann. Statist.*, 42(2):789–817, 2014.
- [134] R. Neal. Markov chain sampling methods for Dirichlet process mixture models. *JCGS*, pages 249–265, 2000.
- [135] R. Neal. Slice sampling. *Ann. Statist.*, 31(3):705–767, 2003.
- [136] R. Neal. *MCMC using Hamiltonian Dynamics*. Handbook of Markov Chain Monte Carlo (S. Brooks, A. Gelman, G. Jones, and X.-L. Meng, eds.), Chapman & Hall / CRC Press, 2010.
- [137] W. Neiswanger, C. Wang, and E. P. Xing. Asymptotically exact, embarrassingly parallel MCMC. In *UAI*, 2014.
- [138] D. Newman, A. Asuncion, P. Smyth, and M. Welling. Distributed inference for latent Dirichlet allocation. In *NIPS*, 2007.
- [139] T. Nickson, M. A. Osborne, S. Reece, and S. J. Roberts. Automated machine learning on big data using stochastic algorithm tuning. *arXiv:1407.7969*, 2014.
- [140] F. Niu, B. Recht, C. Re, and S. J. Wright. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *NIPS*, 2011.
- [141] R. O’Hara and M. J. Sillanpaa. A review of Bayesian variable selection methods: What, how and which. *Bayesian Analysis*, 4(1):85–118, 2009.
- [142] M. Opper. *A Bayesian Approach to Online Learning*. Online Learning in Neural Networks, Cambridge University, 1999.
- [143] J. Paisley, D. M. Blei, and M. I. Jordan. Variational Bayesian inference with stochastic search. In *ICML*, 2012.
- [144] X. Pan, J. Gonzalez, S. Jegelka, T. Broderick, and M. I. Jordan. Optimistic concurrency control for distributed unsupervised learning. In *NIPS*, 2013.

- [145] O. Papaspiliopoulos, G. O. Roberts, and M. Skold. A general framework for the parametrization of hierarchical models. *Statistical Science*, 22(1):59–73, 2007.
- [146] T. Park and G. Casella. The Bayesian Lasso. *JASA*, 103:681–686, 2008.
- [147] S. Patterson and Y. W. Teh. Stochastic gradient Riemannian Langevin dynamics on the probability simplex. In *NIPS*, 2013.
- [148] S. Petrone, J. Rousseau, and C. Scricciolo. Bayes and empirical Bayes: do they merge? *Biometrika*, pages 1–18, 2014.
- [149] N. Pillai and A. Smith. Ergodicity of approximate MCMC chains with applications to large data sets. *arXiv:1405.0182*, 2014.
- [150] J. Pitman. Combinatorial stochastic processes. *Technical Report No. 621. Department of Statistics, UC, Berkeley*, 2002.
- [151] R. Power and J. Li. Piccolo: Building fast, distributed programs with partitioned tables. In *OSDI*, 2010.
- [152] L. R. Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proc. of the IEEE*, 77(2):257–286, 1989.
- [153] R. Ranganath, C. Wang, D. Blei, and E. Xing. An adaptive learning rate for stochastic variational inference. In *ICML*, 2013.
- [154] O. Reichman, M. Jones, and M. Schildhauer. Challenges and opportunities of open data in ecology. *Science*, 331(6018):703–705, 2011.
- [155] D. J. Rezende, S. Mohamed, and D. Wierstra. Stochastic back-propagation and approximate inference in deep generative models. In *ICML*, 2014.
- [156] C. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer, 2005.
- [157] C. P. Robert, J.-M. Cornuet, J.-M. Marin, and N. S. Pillai. Lack of confidence in approximate Bayesian computation model choice. *PNAS*, 108(37):15112–15117, 2011.
- [158] G. Roberts and O. Strame. Langevin diffusions and Metropolis-Hastings algorithms. *Methodology and Computing in Applied Probability*, 4:337–357, 2002.
- [159] F. Ronquist and J. Huelsenbeck. MrBayes: Bayesian inference of phylogenetic trees. *Bioinformatics*, 19(12):1572–1574, 2003.
- [160] A. Roychowdhury, K. Jiang, and B. Kulis. Small-variance asymptotics for hidden Markov models. In *NIPS*, 2013.
- [161] R. Salakhutdinov. Learning deep generative models. *PhD Thesis, University of Toronto*, 2009.
- [162] S. Salihoglu and J. Widom. GPS: A graph processing system. In *SSDBM*, 2013.
- [163] N. Schraudolph, J. Yu, and S. Gunter. A stochastic Quasi-Newton method for online convex optimization. In *AISTATS*, 2007.
- [164] S. Scott, A. Blocker, F. Bonassi, H. Chipman, E. George, and R. McCulloch. Bayes and big data: The consensus Monte Carlo algorithm. *EFA Bayes 250 Workshop*, 16, 2013.
- [165] S. L. Scott. Bayesian methods for hidden Markov models. *JASA*, 97(457):337–351, 2002.
- [166] J. Sethuraman. A constructive definition of dirichlet priors. *Statistica Sinica*, (4):639–650, 1994.
- [167] T. Shi and J. Zhu. Online Bayesian passive-aggressive learning. In *ICML*, 2014.
- [168] A. Smola and S. Narayanamurthy. An architecture for parallel topic models. *VLDB*, 2010.
- [169] J. Snoek, H. Larochelle, and R. P. Adams. Practical Bayesian optimization of machine learning algorithms. In *NIPS*, 2012.
- [170] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *JMLR*, 15:1929–1958, 2014.
- [171] I. Strid. Efficient parallelisation of Metropolis-Hastings algorithms using a prefetching approach. *Computational Statistics and Data Analysis*, 54:2814–2835, 2010.
- [172] M. Suchard, Q. Wang, C. Chan, J. Frelinger, A. Cron, and M. West. Understanding GPU programming for statistical computation: Studies in massively parallel massive mixtures. *JCGS*, 19(2):419–438, 2010.
- [173] M. Tan, I. Tsang, and L. Wang. Towards ultrahigh dimensional feature selection for big data. *JMLR*, (15):1371–1429, 2014.
- [174] M. Tanner and W. Wong. The calculation of posterior distributions by data augmentation. *JASA*, 82(398):528–540, 1987.
- [175] Y. W. Teh, D. Gorur, and Z. Ghahramani. Stick-breaking construction for the Indian buffet process. In *AISTATS*, 2007.
- [176] Y. W. Teh, M. I. Jordan, M. J. Beal, and D. M. Blei. Hierarchical Dirichlet processes. *JASA*, 101(476):1566–1581, 2006.
- [177] Y. W. Teh, A. Thiey, and S. Vollmer. Consistency and fluctuations for stochastic gradient Langevin dynamics. *arXiv:1409.0578*, 2014.
- [178] R. Thibaux and M. I. Jordan. Hierarchical Beta processes and the Indian buffet process. In *AISTATS*, 2007.
- [179] M. Tipping and C. Bishop. Probabilistic principle component analysis. *J. R. Stat. Soc., Series B*, 21(3):611–622, 1999.
- [180] B. M. Turnera and T. V. Zandtb. A tutorial on approximate Bayesian computation. *J. Math. Psychol.*, 56(2):69–85, 2012.
- [181] D. van Dyk and T. Park. Partially collapsed Gibbs samplers: Theory and methods. *JASA*, 103(482):790–796, 2008.
- [182] P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol. Extracting and composing robust features with denoising autoencoders. In *ICML*, 2008.
- [183] M. Wainwright and M. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1–2):1–305, 2008.
- [184] S. G. Walker. Sampling the Dirichlet mixture model with slices. *Commun Stat - Simul and Comput*, 36:45–54, 2007.
- [185] X. Wang and D. B. Dunson. Parallelizing MCMC via Weierstrass sampler. *arXiv:1312.4605*, 2013.
- [186] Y. Wang, X. Zhao, Z. Sun, H. Yan, L. Wang, Z. Jin, L. Wang, Y. Gao, J. Zeng, Q. Yang, et al. Towards topic modeling for big data. *arXiv:1405.4402*, 2014.
- [187] Y. Wang and J. Zhu. Samll variance asymptotics for Dirichlet process mixtures of SVMs. In *AAAI*, 2014.
- [188] Z. Wang, M. Zoghi, F. Hutter, D. Matheson, and N. de Freitas. Bayesian optimization in high dimensions via random embeddings. In *IJCAI*, 2013.
- [189] K. Weinberger, A. Dasgupta, J. Langford, A. Smola, and J. Attenberg. Feature hashing for large scale multitask learning. In *ICML*, 2009.
- [190] M. Welling. Exploiting the statistics of learning and inference. In *NIPS workshop on "Probabilistic Models for Big Data"*, 2013.
- [191] M. Welling and Y. W. Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *ICML*, 2011.
- [192] D. Wilkinson. Parallel Bayesian computation. *Handbook of Parallel Computing and Statistics, Chapter 18*, 2004.
- [193] P. M. Williams. Bayesian conditionalisation and the principle of minimum information. *The British Journal for the Philosophy of Science*, 31(2), 1980.
- [194] S. Williamson, P. Orbanz, and Z. Ghahramani. Dependent Indian buffet processes. In *AISTATS*, 2010.
- [195] S. A. Williamson, A. Dubey, and E. P. Xing. Parallel Markov chain Monte Carlo for nonparametric mixture models. In *ICML*, 2013.
- [196] X.-L. Wu, C. Sun, T. Beissinger, G. Rosa, K. Weigel, N. Gatti, and D. Gianola. Parallel Markov chain Monte Carlo - bridging the gap to high-performance Bayesian computation in animal breeding and genetics. *Genetics Selection Evolution*, 44(1):29–46, 2012.
- [197] R. S. Xin, J. E. Gonzalez, M. J. Franklin, and I. Stoica. Graphx: A resilient distributed graph system on spark. In *Workshop on Graph Data Management Experiences and Systems*, 2013.
- [198] M. Xu, B. Lakshminarayanan, Y. Teh, J. Zhu, and B. Zhang. Distributed Bayesian posterior sampling via moment sharing. In *NIPS*, 2014.
- [199] F. Yan, N. Xu, and A. Qi. Parallel inference for latent Dirichlet allocation on graphics processing units. In *NIPS*, 2009.
- [200] Y. Yu and X.-L. Meng. To center or not to center: That is not the question! an ancillarity sufficiency interweaving strategy (ASIS) for boosting MCMC efficiency. *JCGS*, 20(3):531–570, 2011.
- [201] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica. Spark: cluster computing with working sets. In *Hot Topics in Cloud Computing*, 2010.
- [202] K. Zhai, J. Boyd-Graber, N. Asadi, and M. L. Alkhouja. Mr. LDA: a flexible large scale topic modeling package using variational inference in MapReduce. In *WWW*, 2012.
- [203] A. Zhang, J. Zhu, and B. Zhang. Max-margin infinite hidden Markov models. In *ICML*, 2014.
- [204] X. Zheng, J. K. Kim, Q. Ho, and E. Xing. Model-parallel inference for big topic models. *arXiv:1411.2305*, 2014.
- [205] J. Zhu. Max-margin nonparametric latent feature models for link prediction. In *ICML*, 2012.

- [206] J. Zhu, A. Ahmed, and E. Xing. MedLDA: maximum margin supervised topic models. *JMLR*, (13):2237–2278, 2012.
- [207] J. Zhu, N. Chen, H. Perkins, and B. Zhang. Gibbs max-margin topic models with fast sampling algorithms. In *ICML*, 2013.
- [208] J. Zhu, N. Chen, and E. Xing. Infinite SVM: a Dirichlet process mixture of large-margin kernel machines. In *ICML*, 2011.
- [209] J. Zhu, N. Chen, and E. Xing. Bayesian inference with posterior regularization and applications to infinite latent SVMs. *JMLR*, 15:1799–1847, 2014.
- [210] J. Zhu, X. Zheng, and B. Zhang. Improved Bayesian logistic supervised topic models with data augmentation. In *ACL*, 2013.
- [211] J. Zhu, X. Zheng, L. Zhou, and B. Zhang. Scalable inference in max-margin supervised topic models. In *SIGKDD*, 2013.
- [212] M. A. Zinkevich, M. Weimer, A. Smola, and L. Li. Parallelized stochastic gradient descent. In *NIPS*, 2010.